AD-A154 755 OBJECT IDENTIFICATION AND MANIPULATION USING SILHOUETTE MOMENTS(U) CORNELL UNIV ITHACA NY SCHOOL OF ELECTRICAL ENGINEERING A P REEVES ET AL. APR 85 N00014-83-K-2027 1/1

UNCLASSIFIED F/G 17/5 NL

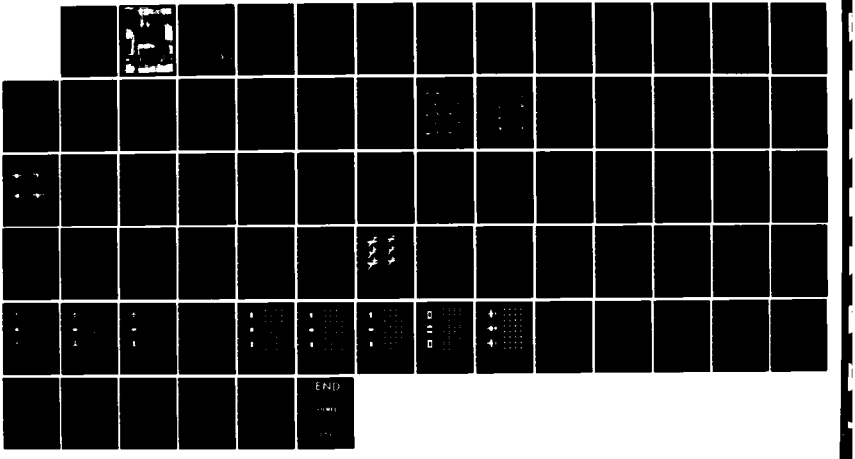

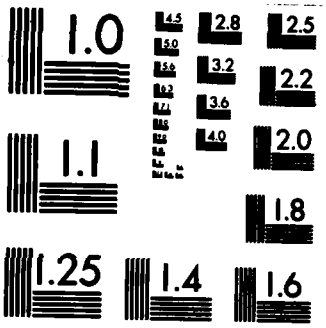

Object Identification and Manipulation Using Silhouette Moments

Anthony P. Reeves
Russell W. Taylor
School of Electrical Engineering
Cornell University
Phillips Hall
Ithaca, New York 14853

April 1985

CORNELL UNIVERSITY

SCHOOL OF ELECTRICAL ENGINEERING
ITHACA, NEW YORK 14853

Object Identification and Manipulation Using Silhouette Moments

Anthony P. Reeves
Russell W. Taylor
School of Electrical Engineering
Cornell University
Phillips Hall
Ithaca, New York 14853
April 1985

Final report for NRL Contract # N00014-83-K-207

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered) A154 755

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Study of Autonomous Image Processing Algorithms: Object Identification and Manipulation Using Silhouette Moments		5. TYPE OF REPORT & PERIOD COVERED Final report July 1983-December 1984 6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) A.P. Reeves R.W. Taylor		8. CONTRACT OR GRANT NUMBER(s) N00014-83-K-2027
9. PERFORMING ORGANIZATION NAME AND ADDRESS School of Electrical Engineering Cornell University Ithaca, NY 14853		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
11. CONTROLLING OFFICE NAME AND ADDRESS		12. REPORT DATE April 1985 13. NUMBER OF PAGES
14. MONITORING AGENCY NAME & ADDRESS(if different from Controlling Office)		15. SECURITY CLASS. (of this report) 15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)
Approved for public release: distribution is unlimited

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

18. SUPPLEMENTARY NOTES

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)
Shape Analysis, moments, Fourier descriptors, pattern recognition, noise reduction, FLIR imagery, silhouette imagery, inverse moment transform, Legendre moments, image averaging, moment subset classification, six airplane experiment, aspect ratio normalization, variance balancing.

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)
The theory of silhouette moments are presented along with results for classifying clean and noisy images. Noise reduction by image averaging is examined comparatively in image and moment spaces. Results are presented for classifications using subsets of the complete moment set. An inverse moment transform procedure is developed and demonstrated. and [illegible] include

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-LF-014-6601

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

Abstract

This report considers object identification and manipulation using silhouette moments. The fundamental theory behind moments for image analysis is presented, and comparative object identification results are given for silhouette moments and a number of other object description techniques. Also discussed are the techniques of aspect ratio normalization, and variance balancing, which greatly enhance the classification success. Both clean and noisy synthetic data are considered, with silhouette moments the most successful and economical technique in both cases. A further consideration for noisy data is the reduction of the noise level; to accomplish this, multi-image averaging is examined comparatively in the moment and image domain. It is shown that averaging in the moment domain is particularly effective in situations where the location of an object within a scene is uncertain. The optimization of the elements contained in the feature vectors is also of importance to achieve accurate and efficient object identification. Feature vectors composed of subsets of complete silhouette moment sets are examined, and it is found that odd order moments in general, and order two moments, are critical to classification success. A final consideration is the inverse transform from moment domain to image domain. An iterative technique to accomplish this is presented, along with actual results.

TABLE OF CONTENTS

1. Introduction

The suitability of the method of moments for object identification has been demonstrated in the literature repeatedly. The most fundamental moment based approach is to use only the silhouette of an object to produce *silhouette moments*. This report will consider a number of issues related to object recognition using silhouette moment representations.

The first section contains an examination of the fundamentals of moment theory and object normalization. A comparison of classification results for the techniques of silhouette moments, moment invariants, Fourier descriptors, and several other methods are presented. This discussion provides the basis for the consideration of moments as an object identification tool. A system for evaluation of experimental object identification success is also presented, and then further refined in section 4.

The comparison experiments presented in section 1 were funded by the National Science Foundation and the Army. They are presented in more detail in [5], which is also included as appendix A. These results were included to maintain the completeness and continuity of this report.

Given the possibility of obtaining multiple images of a noisy object, an examination of noise reduction by image averaging in the moment domain is of interest. In section 3 a technique for averaging image sequences in the moment domain is presented, and a comparison of averaging in the greylevel and moment domains is made. In cases where the object's precise position within the image is not known with certainty from frame to frame, it will be shown that moment averaging can be very effective for noise reduction.

As an extension of the object identification work, it is of interest to know the relative importance of different moment values to object recognition. In section 4 an examination is made of object identification success using various subsets of the standard complete moment set. This analysis is performed with various types of moment normalization, and using clean and noisy objects.

Section 5 presents an inverse moment transformation to convert from a set of silhouette moments back to a silhouette image. The performance of this technique is discussed, and a number of sample reconstructions are presented.

2. Fundamental Moment Theory and Experimentation

The properties of moments and Fourier descriptors have been examined by the use of an experiment involving six airplane shapes viewed at random angles in three dimensions. Fourier descriptors as used by Wallace and Wintz [1] have been compared with our own scheme [2] for normalized moments. Moment invariants as used by Dudani et. al. [3]. have also been considered. Many different normalized moment forms have been proposed [2] including Legendre moments and variants of moment invariants [4] In this paper the performances of several normalization techniques were compared.

Most of the previous work with moments, such as that by Dudani et al [3]. has centered on moment invariants; we have concentrated our efforts on normalized conventional moments which we consider to be less sensitive to noise. Section 2.2 of this report outlines the theory of conventional moment normalization. In section 2.3 a new normalization technique, aspect ratio normalization, is described. We have compared the performance of moments with Fourier descriptors; section 2.4 briefly outlines the method we used for computing Fourier descriptors. In section 2.5 details of our classification procedures are given. The results for experiments are given in section 2.6.

The first six airplane experiment was conducted by Dudani et. al.[1] using moment invariants as the feature vectors. A library of feature vectors for different views of each aircraft were constructed and then an unknown was identified by the best match into these libraries. The input image data was obtained with a TV camera viewing a plastic model of each airplane. Wallace and Wintz [2] conducted a similar type of experiment to investigate the properties of Fourier descriptors as feature vectors. In this case the airplane images, for both libraries and unknowns, were generated by the computer from a 3-dimensional graphics model. This procedure is more convenient than Dudani's in that data from any angle or resolution can be computed independent of the imaging system defects.

The experiments described in this section were based on the software used by Wallace and Wintz. The same six airplanes were used and also the same method was used to generate

the Fourier descriptor feature vectors. However, most other parameters and the feature vector classification scheme was significantly different.

In our experiments the minimum Euclidean distance in the feature space of the unprocessed feature vectors was used to determine distances. A feature vector weighting scheme was also investigated which set the expected value for each element of the total library to one. This produced higher classification success rates for moment feature vectors.

The comparison experiments presented in this section were funded by the National Science Foundation and the Army, and are presented here for completeness. A more detailed discussion of these experiments is contained in [5] which is included as Appendix A.

2.1. The Six Airplane Experiment

A standardized six airplane classification task was used for all experiments. It is loosely based on an experiment designed and developed by Wallace and Wintz [1]. Their experiment was, in turn, based on the experiment conducted by Dudani et.al. [3].

In the six airplane experiment the six airplanes were described by a short list of parameters. These parameters are used to generate a representation of the three dimensional surface of an airplane by a set of three and four sided polygons. A library of feature vector sets of two dimensional views for each airplane was generated. For the experiments in this section, library samples were usually taken over a uniform grid of viewing angles. Each library entry was normalized with respect to size, location, and in-plane rotation.

A set of random views of the six airplanes was used to generate a test set of feature vectors. Each test view was matched by means of nearest Euclidean neighbor in feature space to its closest entry in the library. The class of this entry was then used to identify the test airplane. The performance of the shape analysis scheme was indicated by the number of correct classifications.

In our tests we increased the library size from 143 views covering one hemisphere, as used by Wallace and Wintz, to 500 equally spaced views covering the whole sphere of possi-

ble viewing directions. All the results were computed with 500 library views for each of the six airplanes and the standard 50 unknowns as used by Wallace and Wintz. All image resolutions were 128 x 128 for both libraries and unknowns.

2.2. Moment Feature Vectors

The conventional definition of the two-dimensional moment of order (p+q) of a function f(x,y) is

$$M_{pq} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x^p y^q f(x,y)\, dx\, dy$$

$$p,q = 0,1,2, \cdots \tag{2.1}$$

If f(x,y) is piecewise continuous and has nonzero values only in the finite part of the x-y plane, then moments of all orders exist. Furthermore, the double moment sequence $\{M_{pq}\}$ is uniquely determined by f(x,y); and conversely f(x,y) is uniquely determined by $\{M_{pq}\}$ [6].

A set of moment values may be used to represent a segment of an image. Operations such as rotation, translation and scale change may more easily be achieved in the moment domain than the original pixel domain. Furthermore, a truncated set of moments may offer a more convenient and economical representation of the essential shape characteristics of an image segment than a pixel based representation.

There has been sporadic interest in the use of image segment moments for shape analysis. One of the earliest relevant theoretic studies of moments was presented by M-K Hu in 1962 [6]. He described how fundamental transformations such as translation, rotation, and scale change could be achieved with the moment image representation. Rotation invariance proved to be the most difficult to achieve. Hu proposed two methods for rotation invariance; aligning the moments with the principle axis and computing rotation invariant combinations of moments called moment invariants.

A second important theoretical work on moments was R. M. Teague's 1980 paper [4]. Teague introduced the concept of orthogonal moment sets and showed how the concept of Hu's moment invariants may be extended. He also showed how an approximate inverse

moment transform could be defined. Hu defined seven moment invariants of order three and less. He demonstrated the method of invariants for a character recognition application in which two moment invariants of order two were used as a feature vector.

Recent work [7] has shown that both the moments of the binary silhouette and the grey-valued image segment may be used together for classification. Furthermore, the two dimensional moments may be used to obtain moments of projections through the image segment which may, in turn, be used in classification schemes.

2.2.1. Transformations of Conventional Moments

The operations of scale change, translation and rotation of an image segment are easily defined in terms of moments

(a) Scale change

The moments of f(x,y) after a scale change by a factor λ are defined by

$$M'_{pq} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x^p y^q f(x/\lambda, y/\lambda) dx\ dy \tag{2.2}$$

this may be expressed in terms of the original moments of f(x,y) by

$$M'_{pq} = \lambda^{2+p+q} M_{pq} \tag{2.3}$$

(b) Translation

The moments of f(x,y) translated by an amount (a,b) are defined by

$$M'_{pq} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (x+a)^p (y+b)^q f(x,y) dx\ dy \tag{2.4}$$

this may be expressed in terms of the original moments of f(x,y) by

$$M'_{pq} = \sum_{r=0}^{p} \sum_{s=0}^{q} \binom{p}{r} \binom{q}{s} a^{p-r} b^{q-s} M_{rs} \tag{2.5}$$

(c) Rotation

The moments of f(x,y) after a rotation by an angle θ about the origin are defined by

$$M'_{pq} = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} (x\cos\theta + y\sin\theta)^p (y\cos\theta - x\sin\theta)^q f(x,y) dx\ dy \tag{2.6}$$

this may be expressed in terms of the original moments of f(x,y) by

$$M'_{pq} = \sum_{r=0}^{p} \sum_{s=0}^{q} \binom{p}{r} \binom{q}{s} (\cos\theta)^{p-r+s} (\sin\theta)^{q+r-s} M_{p+q-r-s,r+s} \qquad (2.7)$$

The rotation of an image segment about an arbitrary point may be achieved in three steps. First translate the moments so that the point of rotation is at the origin using (2.5), then rotate the moments using (2.7) and, finally translate the moments such that the point of rotation is returned to its original position.

(d) Reflection

A reflection about the y axis results in a change of sign of all moments which depend upon an odd power of y, i.e.,

$$\begin{aligned} M'_{pq} &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} x^p \, y^q \, f(x,-y)\, dx\, dy \\ &= (-1)^q M_{pq} \end{aligned} \qquad (2.8)$$

A reflection about an arbitrary axis may be achieved in three steps. First, the moments are rotated so that the reflection axis is aligned with the y axis. Second, the moments are reflected about they axis by using the transform specified by (2.8). Finally, the reflected moments must be rotated to restore the reflection axis to its original orientation.

2.2.2. Normalized Moment Sets

For image analysis purposes an image segment is represented by a truncated set of moments. From inspection of equations (2.3), (2.5), (2.7) and (2.8) in the previous section we see that a transformed moment value does not depend upon any moment of a higher order. A *complete moment set* (CMS) of order n is defined to be a set of all moments of order n and lower. Operations which correspond to translation, rotation and scale change are closed with respect to a CMS of f(x,y).

A CMS of order n may be conveniently expressed in a triangular matrix form as shown below:

$$\begin{matrix} M_{00} & M_{01} & M_{02} & \cdots & M_{0n} \\ M_{10} & & & & \\ \cdot & & \cdot & & \\ \cdot & & \cdot & & \\ \cdot & & \cdot & & \\ M_{n0} & & & & \end{matrix}$$

A CMS contains (n+1)(n+2)/2 elements.

Scale normalization may be achieved by setting M_{00} to 1. For a binary image or a grey level image with a mean grey level of 1 this corresponds to standardizing the area of the image segment to 1. The scale factor to achieve this normalization is

$$\lambda_N = 1/\sqrt{M_{00}} \tag{2.9}$$

A slightly different scheme is used by Dudani et al. The radius of gyration is normalized to 1 by using the following scale factor

$$\lambda_R = 1/\sqrt{\mu_{02} + \mu_{20}} \tag{2.10}$$

Where μ_{02} and μ_{20} are central moments as defined by equation (2.13). It is expected that the results of the two schemes would be very similar.

Translation normalization is usually achieved by translating the origin to the center of gravity of the image; moments with their origin at the center of gravity of the image are called central moments. The position of the center of gravity $(\bar{x},\bar{y})$ may be determined from non central moments by

$$\bar{x} = M_{10}/M_{00} \tag{2.11}$$

$$\bar{y} = M_{01}/M_{00} \tag{2.12}$$

The central moments μ_{pq} can then be computed from (2.10) substituting $a = -\bar{x}$ and $b = \bar{y}$ i.e.,

$$\mu_{pq} = \sum_{r=0}^{p}\sum_{s=0}^{q} \binom{p}{r}\binom{q}{s} (-\bar{x})^{p-r}(-\bar{y})^{q-s} M_{rs} \tag{2.13}$$

Two methods of rotation normalization have been proposed by Hu [6]. In the first method the principal axis of the image are determined and the moments are transformed to align with these axes. The main problem with this method is to uniquely determine the principal axis. The second method is to generate simple combinations of moments called moment

invariants which Hu derived from algebraic invariants. In this case it is not necessary to determine the principal axis but it is difficult to define further operations on sets of moment invariants.

The principal axis moments are obtained by rotating the axis of the central moments until μ_{11} is zero. The angle θ, from the original axis to the principal axis is defined by

$$\tan 2\theta = \frac{2\mu_{11}}{\mu_{20} - \mu_{02}} \tag{2.14}$$

The angle theta obtained with (2.14) may be with respect to either the major principal axis or the minor principal axis. One way to determine a unique orientation of the principal axis is to set the additional constraints that $\mu'_{20} > \mu'_{02}$ and $\mu'_{30} > 0$. The correct rotation angle will be $\theta + n\pi/2$ where n is chosen to satisfy those constraints. Problems may occur when an image is n-fold symmetric since there are multiple possible sets of principle axis.

For some applications reflection normalization may be useful. In this case an image segment and its mirror image would be represented by the same normalized moment set. For moments with the above rotation normalization a reflected image would be reflected about the y axis when compared with a non-reflected image. Therefore, an additional constraint such as $\mu'_{30} > 0$ may be used to ensure reflection normalization. If after rotation normalization $\mu'_{30} < 0$ then a reflection transformation as specified by (2.8) would be necessary.

2.2.3. Standard Moments

A set of standard moments is a CMS which has been normalized with respect to scale, translation, rotation and optionally reflection. The low order standard moments have the following values:

$M_{00} = 1$ {area normalized to 1}

$M_{10} = 0$ {central moments, taken about

$M_{01} = 0$ the center of gravity}

$M_{11} = 0$ {rotation normalization; moments are taken about the principal axis}

For rotation normalization we also have

$$M_{20} \geq M_{02}$$

$$M_{30} \geq 0$$

If reflection normalization is selected

$$M_{03} \geq 0$$

The standard moment representation of a segment is very convenient for shape analysis and may be compared directly with other standard moment sets for shape similarity.

2.3. Aspect Ratio Normalization

A feature of the standard moments, as described in the previous section, is that for blob-like objects the magnitude of the moments will decrease as the order increases. The condition for diminishing magnitudes is that the object be confined within a 2x2 square centered at the origin. Recall that M_{00}, the object area, is 1, therefore all other moments will be less than 1. If any part of the object extends beyond this square region, then some moments will become increasingly large and even for low orders may be much larger than 1. A significant number of the airplane library images extend beyond the 2x2 square.

In order to improve the behavior of feature vectors, an Aspect ratio normalization was developed. This is a projective transform of an image of the object which converts the ellipsoid of inertia of the object to a circle, i.e., $M_{20} = M_{02}$, while maintaining $M_{00} = 1$. There are still shapes which would extend beyond the 2x2 square even after this normalization; however, all of the airplane images were contained in this square.

A change in scale by a factor of δ in the x direction and γ in the y direction is given by:

$$M'_{pq} = \delta^{1+p}\, \gamma^{1+q}\, M_{pq} \tag{2.15}$$

For aspect normalization the area remains 1, i.e.,

$$M'_{00} = M_{00} = 1 \tag{2.16}$$

From (2.15)

$$M'_{00} = \delta\gamma M_{pq}$$

$$\delta = 1/\gamma$$

For aspect normalization

$$M'_{20} = M'_{02}$$

From (2.15)

$$\delta^3\gamma M_{20} = \delta\gamma^3 M_{02} \tag{2.17}$$

$$\delta = (M_{02}/M_{20})^{0.25}$$

The aspect transformation is achieved with (2.15) using δ, the aspect ratio as defined in (2.17), and $\gamma = 1/\delta$.

This transformation is applied to standard moments. For the transformed moments, $M_{20} = M_{02}$, therefore we discard one of them. The aspect ratio δ is an important shape feature and is used as the first or most significant feature element. Its value is in the range $0 < \delta \leq 1$.

2.4. Fourier Descriptors

A common set of features used to recognize a silhouette or contour is the Fourier descriptor of the contour [8]. The method used in our experiments is fully described in [1]. A short description of the process is included here.

Consider a closed contour C in the complex plane. Trace it once in the counter-clockwise direction with uniform velocity v, obtaining the complex function z(t) with parameter t. Choose v so that the time T required to traverse the contour is 2π. Traversing the contour more than once yields a periodic function, which may be expanded in a convergent Fourier series. We will define a Fourier descriptor of C to be the complex Fourier series expansion of z(t).

$$z(t) = \sum_{n=-\infty}^{\infty} A(n)\, exp(jnt)$$

where

$$A(n) = \frac{1}{2}\pi \int_0^{2\pi} z(t)\, exp(-jnt)\, dt$$

This FD depends on both C and the starting point of z(t). In practice, C is taken from a digitized image, and thus z(t) is not available as a continuous function. If z(k) is a uniformly sampled version of z(t) of dimension N, the discrete Fourier transform directly gives us the N lowest frequency coefficients A(n).

The computation of the FD is fairly straightforward. The original sampled contour is assumed to be given in the form of an eight-neighbor chain code. This chain code representation is converted into a sequence of x-y coordinates. The piecewise linear contour defined by the original chain code can be smoothed as described in [1] and the resulting sequence of x-y coordinates is taken to accurately represent the original contour. The length of the contour is computed, and the contour is resampled at a spacing chosen to make the total number of samples a power of two. The FD is computed by simply taking the FFT of this sequence. Alternatively the Fourier coefficients can be calculated using a DFT approach which uses the piecewise linear nature of the chain code to speed the calculation.

The frequency domain operations which affect the position, size, orientation, and starting point of the contour follow directly from properties of the DFT. Normalization of an unknown contour with respect to these variables is required before comparison with a library of known shapes is possible. Let the input FD be A(i), $-N/2 + 1 \leq i \leq N/2$. Let the coefficient of second greatest magnitude be A(k) ($k \neq 0$). Let u and v denote the phases of A(1) and A(k), respectively.

1. Set A(0) = 0. (This removes translation effects).
2. Divide the A(i) by |A(i)|. (This normalizes the scale).
3. Multiply the A(i) by $exp([(i-k)u + (1-i)v]/(k-1))$. This simultaneous application of the rotation and starting point shift operations sets the phase of both A(1) and A(k) to zero. It is not unique; there are |k-1|-1 other possible normalizations which also do this.

4. If k=2, normalization is complete. Otherwise proceed to step 5.

5. Compute the sum $\sum_{i=1}^{N-1} Re[A(i)] \mid Re[A(i)] \mid$, and save. We will use the normalization which maximizes this sum.

6. If step 5 has been executed for all |k-1| different normalizations, go to step 7. Otherwise multiply the A(i) by $exp(j(i-1)2\pi/\mid k-1\mid$ and go to step 5.

7. Find the largest ambiguity-resolving criterion from step 5 among the |k-1| values previously saved. The associated normalization defines the NFD.

The normalized Fourier descriptor thus obtained consists of coefficients representing harmonics of the contour.

2.5. Classification Procedures

Each airplane, considered from all possible viewing angles, forms a continuous surface in feature space. The novel classification problem is to find the closest airplane surface to an unknown airplane, which is represented by a single point in feature space. Classical feature vector conditioning procedures, such as the Eigen value transformation as used by Wallace and Wintz, are not necessarily appropriate in this case since the within-class feature vector variation (for all different possible viewing angles) is, in general, less than the between-class variation for a given viewing angle.

We have used two classification criterias. The first is simply minimum Euclidean distance in feature space. Both standard moments and normalized Fourier descriptors have feature elements which decrease in magnitude with order or frequency. Therefore, more weight is placed on the lower order or frequency elements which are expected to be more robust with respect to noise.

The second criteria is minimum Euclidean distance after "variance balancing" the feature vectors. Variance balancing is a weighting technique to make the effect of the feature elements more equal. The variance of each feature element for all feature vectors in the total library for all six airplanes is first computed. Then all feature elements are weighted by the

inverse of the standard deviation computed for the library. Therefore, if all views of the airplane are equally likely, the expected variance of a weighted feature vector is one.

2.6. Experiments and Results

2.6.1. The Basic Experiment

The following parameters were used in the basic experiment reported in this section.

(1) The six airplanes were: a B57, an F104, an F105, a MiG, a Mirage and a Phantom. These were the same planes used by Dudani and by Wallace and Wintz.

(2) The library views consisted of 500 views uniformly sampled over the total sphere of viewing angles. In Dudani's experiment 550 views of each airplane were taken at 5 degree intervals in a 140 by 90 degree sector. The library views used by Wallace and Wintz consisted of 143 entries covering one hemisphere in a 11 x 13 grid.

(3) The test set consisted of the same 50 random views of each airplane for a total set size of 300. The viewing angles used were identical to those of Wallace and Wintz.

(4) Library entries were computed from two-dimensional views digitized on a 128 x 128 pixel matrix. Test set entries were also computed from a 128 x 128 pixel representation.

Representative views of each of the six airplanes with a 128 x 128 pixel resolution are shown in figure 2.1. Typical random views of noisy versions of these airplanes are shown in figure 2.2.

2.6.2. Noisy Data Generation

Considerable care was taken in the generation of the noisy data to give realistic results. The following operations were applied to each clean test image.

1. The silhouette image was converted to a grey level image in which each object point is 64 greater than the background.

2. A 2 x 2 local mean filter operation is performed which blurs the edges of the grey level image.

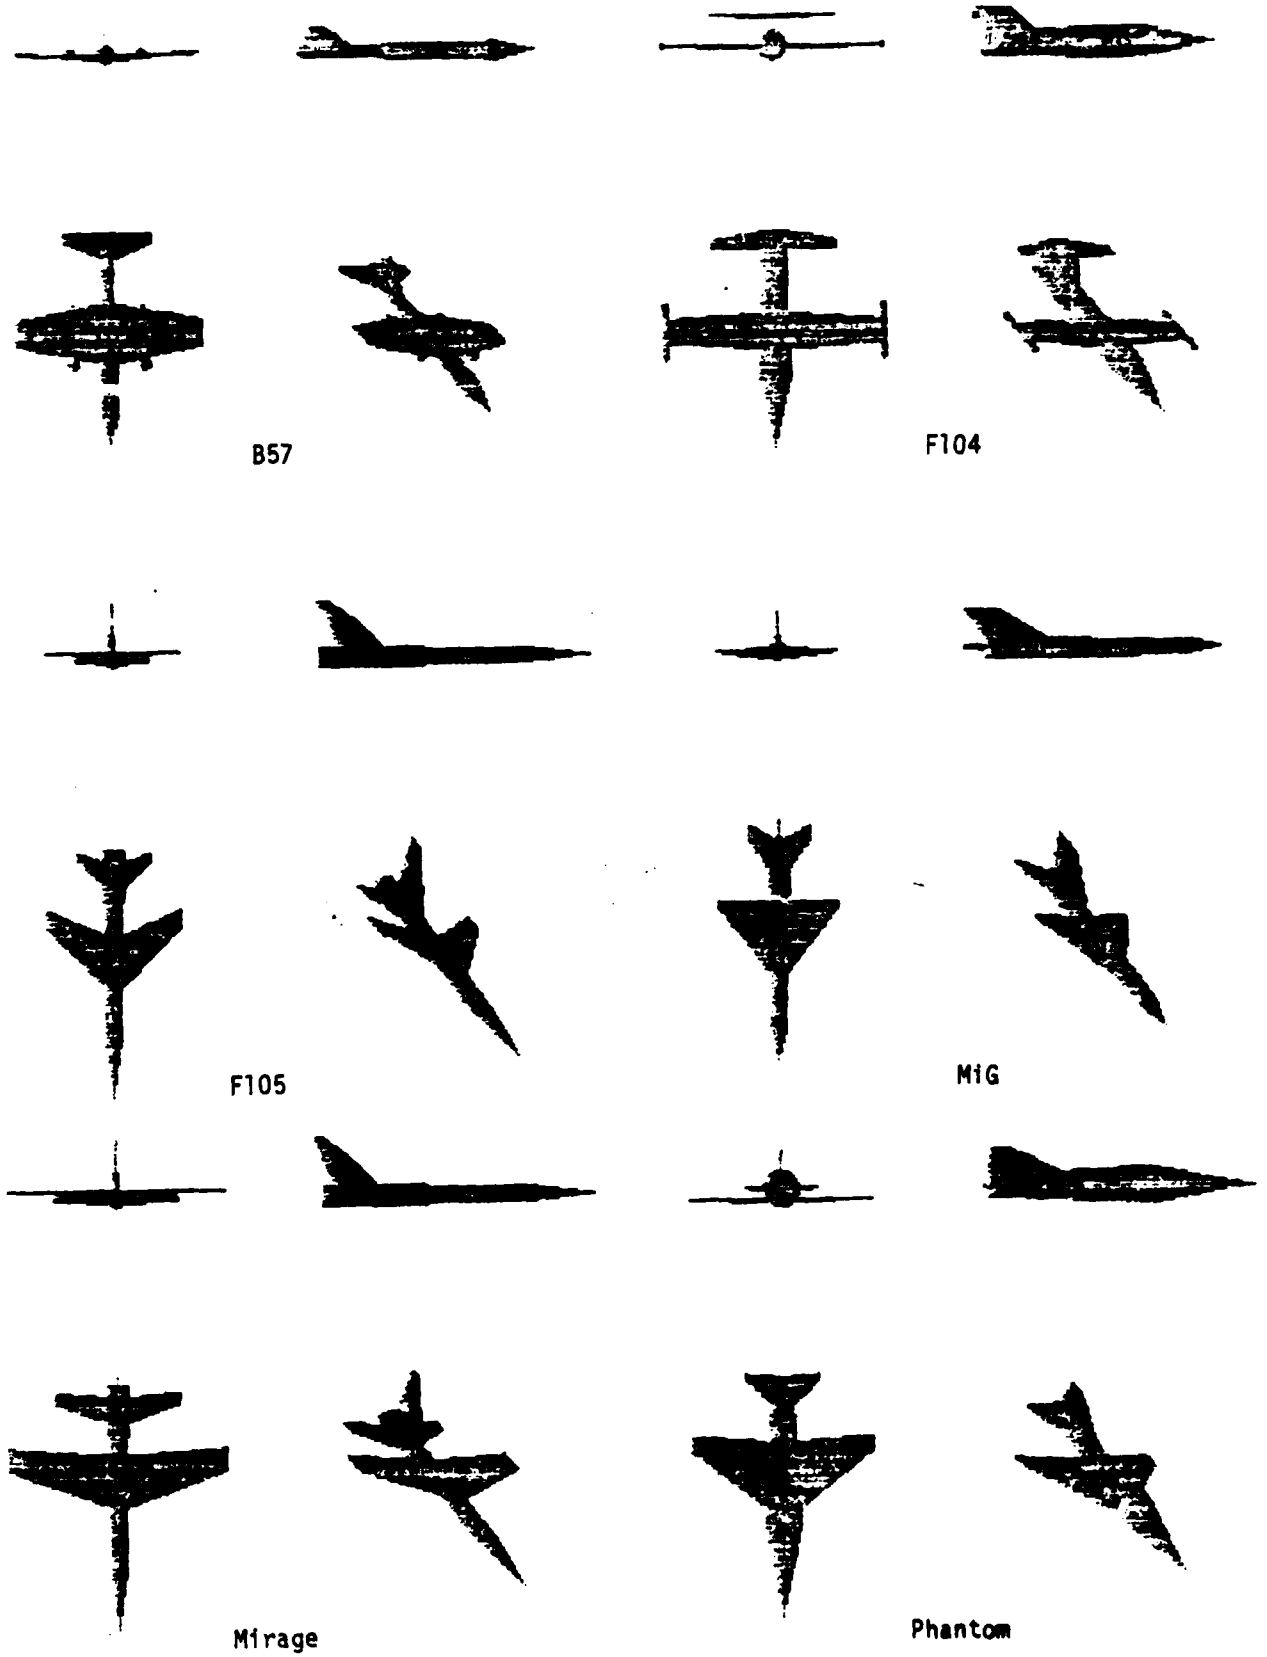

Figure 2.1.
Representative views of the six airplanes used in the moment based classification tests.

Figure 2.2.
Sample of random views of noisy versions of the six airplanes.

3. Gaussian noise is added with a mean of zero and a standard deviation of 30.

4. A second 2 x 2 local mean filter operation is performed.

5. A thresholding operation is performed at the optimal grey level for the image specified in step 1.

6. A binary filtering algorithm is applied which selects only the main connected region of ones (object points) in the image.

2.6.3. Results

Fourier descriptors exhibit lower classification success than silhouette moments alone for comparably sized feature vectors. For small size feature vectors, variance balancing is beneficial but for larger numbers of elements the results are worse. The results for clean data are slightly better than those obtained by Wallace and Wintz; this could be due to the denser library used.

For conventional moments and aspect ratio normalized moments variance balancing gives a significant improvement classification success. The best results are obtained with variance balanced aspect ratio normalized moments. This performance is effectively achieved with moments of order 3; using more moment values gives slight or no improvement. These results are consistent for both clean and noisy data sets.

Figure 2.3 shows the performance of the various types of moments, and Fourier descriptors, on clean data. More extensive results are given in [8]. Combining silhouette and boundary data did give a slight improvement over the silhouette data alone. Moments based on Legendre polynomials have been proposed by Teague [4]. While the orthogonal property might be expected to lead to better feature vector representation, our results indicate that they performed worse than the silhouette moments from which they were derived. Boundary moments by themselves performed much worse then silhouette moments.

Also shown in figure 2.3 are results for moment invariants which are worse than any other technique tested. These were computed in a similar way to Dudani [3] except for the

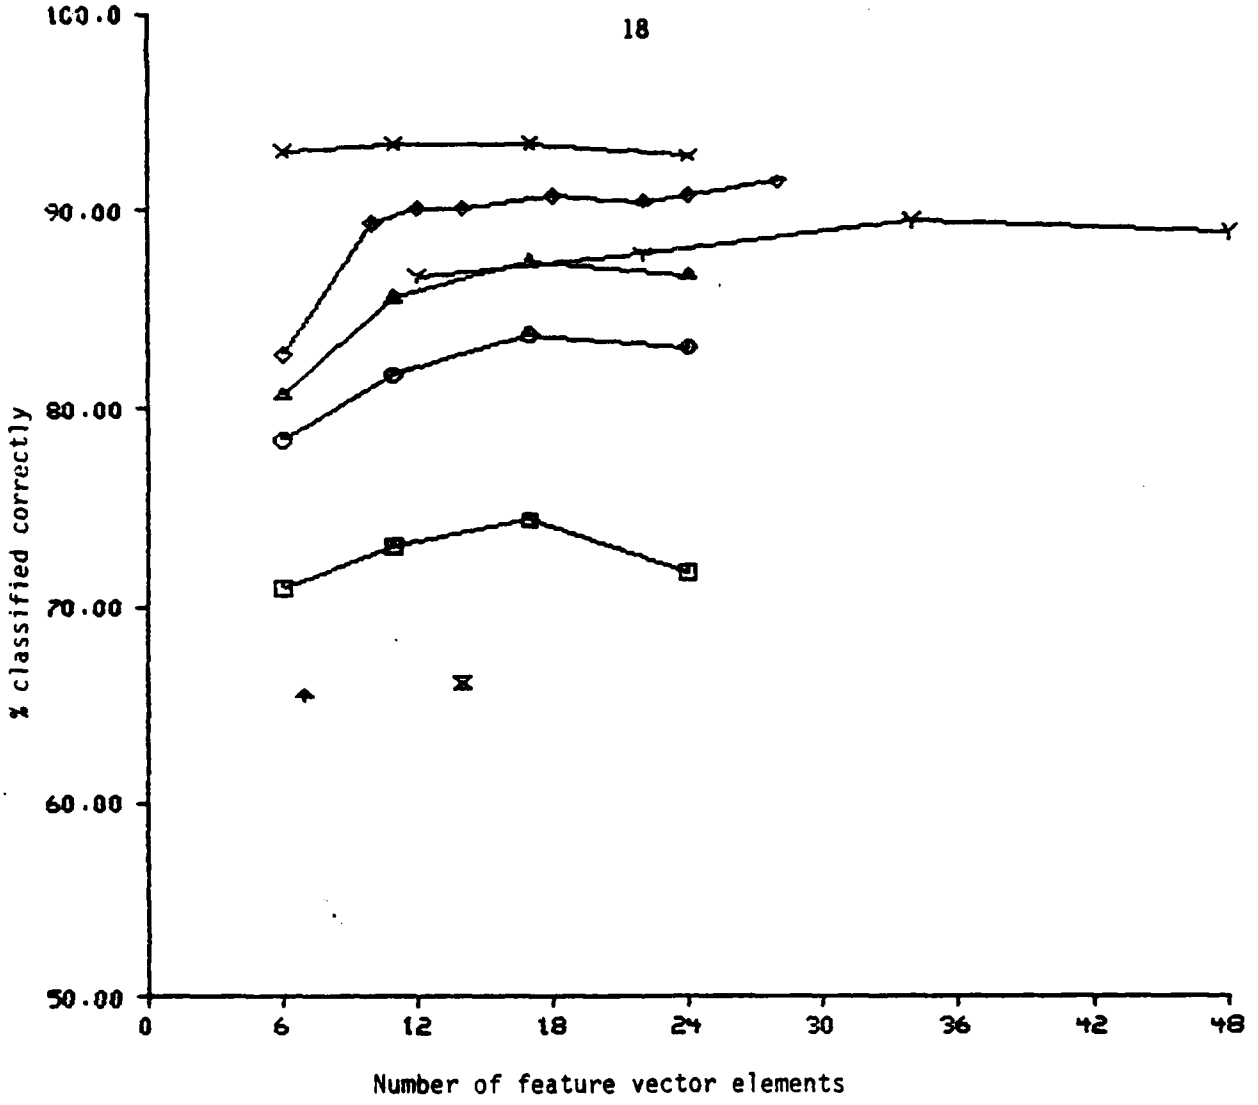

Figure 2.3
Graph showing results of classification of clean data. All results are variance balanced.

first feature element value. Dudani assumed that the distance to the object was known; therefore the absolute scale of the object was also known and is embodied in the first feature vector element. In our studies we have assumed that the distance to the object is not known. In our moment invariants the radius of gyration (the first feature element) is size normalized by the perceived area of the object M_{00} while Dudani ignores this area moment value.

2.7. Summary of Results

The following summarizes the best results which have been obtained on clean and noisy test data for the two most promising techniques, aspect ratio normalized silhouette moments and Fourier descriptors:

Classification Success		
No. of Feature Vector Elements	Percentage Correctly Classified	
	Standard Moments	Fourier Descriptors
6	93.0 (85.7)	77.7 (75.0)
11	93.3 (89.7)	90.7 (82.3)
17	93.3 (89.7)	90.3 (85.0)
24	92.7 (91.0)	91.7 (85.6)

The figures in parenthesis are for a noisy set of unknowns. For moment invariants of the silhouette (7 feature vector elements) the highest recognition rate was 68.0% for clean data and 69.3% for noisy data. Adding boundary moment invariants gave worse results. The above results indicate that there is very little difference between the performance of Fourier descriptors and standard moments for this task; both do very well. Moment invariants did considerably worse than the other techniques. The method used for normalization significantly affects the results. In general, the addition of boundary moments did not improve the classification results for a given number of feature vectors and boundary moments by themselves were not very good.

3. Analysis and Modelling of the BQM-34 Drone Imagery

The FLIR images of the BQM-34 drone plane provide a database of high noise representations, exemplifying the type of image corruption introduced by atmospheric inhomogeneities. The analysis and understanding of this type of data is necessary to obtain robust image identification procedures. In particular we would like to examine the effects of such noise on image classification by the method of moments.

One method for reducing additive noise is the averaging of multiple distinct images. This report examined the correspondence of varying size multi-image group averages to an *ideal* model. By making the comparison in the moment domain for both the averaging of moment sets, and of greylevel images, we compared these two noise removal approaches from the standpoint of producing feature vectors suitable for image identification.

For realistically modelled image aquisition, the ease of location and orientation normalization of moments makes the method potentially very attractive for the reduction of noise effects by averaging. In particular we examined the situation where there was uncertainty about the object's location within an image frame. This scenario would be typical of most situations where one or both of the observer and the observed object are moving.

3.1. Experimental Methods

3.1.1. The Geometric Model

One of the first concerns was to produce an idealized version of the BQM-34 drone plane to create both a reference for comparisons and a base for noise simulations. A polygonal representation of the drone was produced based on published measurements. This representation was then rotated until a visual best match to real drone data was found. This orientation was then used for all experiments. Care was also taken to match the resolution of the real data as closely as possible.

Observation of the FLIR drone images in comparison to the published specifications indicated some discrepancies. Most notably, the pronounced tail cone section was not visible in the real data. This might be an artifact of the range imaging process, but in any case it is a source of variation between the model and real data moments. It is important to realize that a perfect match is possible between the noisy model and our ideal, since the noisy images are based on the ideal. No such perfect match is possible due to averaging for the real data, since the ideal for the real data is not identical to the model ideal. Figure 3.1 shows some representative real and model images.

3.1.2. Noise and Preprocessing

The drone FLIR data contained substantial real atmospheric and measurement noise. These drone images were preprocessed prior to taking the moments of an individual image by a 3 by 3 mean filtering. This filtering removed stray background pixels that would otherwise be misclassed as object when segmented from the background. For the multi-image averages the mean filtering was found to be unnecessary.

The artificial drone images received a slightly more complex treatment. To produce the noisy images, an ideal silhouette of the model drone was created with an intensity value of 128. Gaussian noise with a standard deviation of 30 was then added to this model to produce the noisy images.

These artificial noisy images were then mean filtered, thresholded, and processed to remove noise pixels that were not 8 connected to the main body of the object. The latter processing step removed spurious noise spots in the image, as did the mean filtering used on the real drone data. However, a higher degree of object distortion was maintained by the non-mean filtering technique, which allowed the noisy model data to more accurately simulate the real drone data.

The preprocessing for noise reduction was intentionally kept simple, to avoid obscuring the effects of the greylevel averaging. The fundamental interest of this work concerned the quantification of the effects of turbulence and measurement noise, and the reduction of these

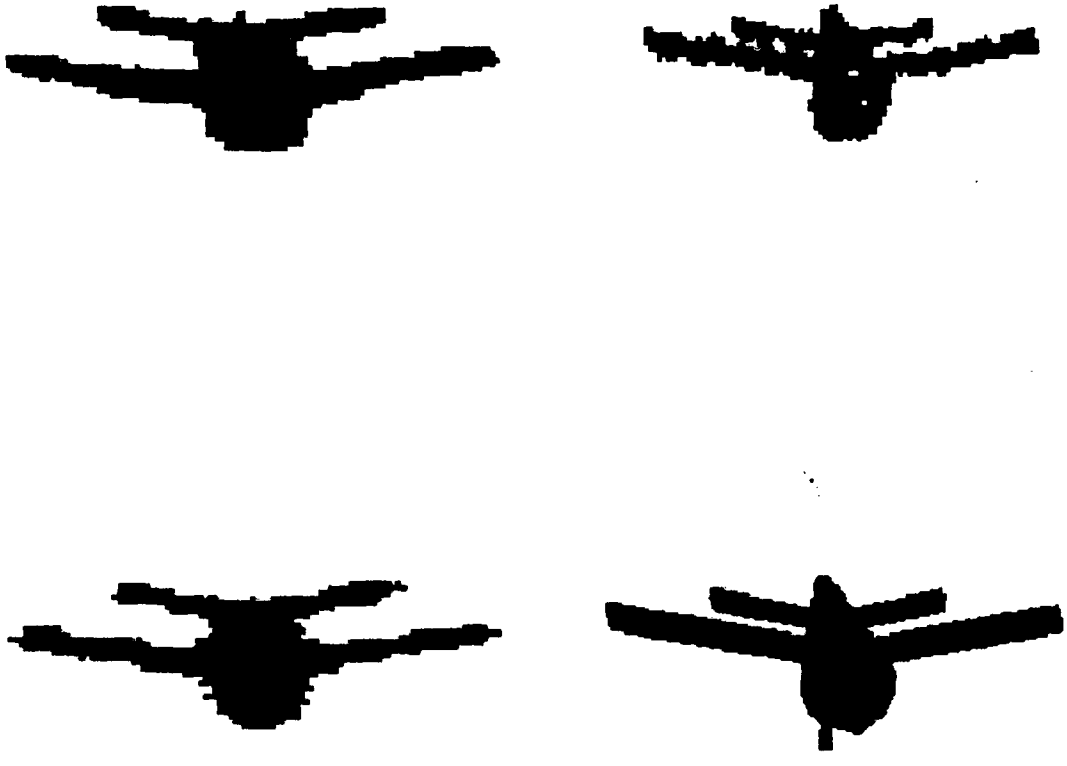

Figure 3.1.
Samples of the real and model BQM-34 drone data. Top left: real drone image after noise preprocessing and thresholding. Top right: noisy model drone image after noise preprocessing and thresholding. Lower left: threaholded version of the average of all 40 real drone images. Lower right: the ideal model image without noise.

effects afforded by image averaging. In particular, the extra processing of the artificial images was performed solely to produce a reasonable simulation of the real data.

3.1.3. Greylevel Versus Moment Domain Averaging

The basic concern of this experimentation is to determine the relative merits of averaging images in the moment domain versus averaging in the image domain. The final intent is to produce silhouette moment-based feature vectors with reduced noise effects. A secondary question, for either of these two techniques, is what effect does the number of images averaged have on the convergence of the averaged result to the ideal.

The technique used to address both these concerns involved the partitioning of data sets of 40 images into groups containing 1, 5, 10, 15, 20, and 40 distinct and different images (specifically 40, 8, 4, 2, 2, and 1 groups, respectively). The members of these groups were then averaged in both the greylevel and the image domain, and the silhouette moment set associated with each of these two averagings was then compared to the ideal silhouette moment set. The figure of merit for the comparison was the Euclidean distance between the moment sets. All the Euclidean distances for a given group size and averaging technique were then averaged together to come up with a representative distance figure for groups of that particular size.

Two types of greylevel averaging are possible. The first is to actually add up a series of greylevel values for each pixel point and then determine the actual pixel average. This averaged image is then thresholded to produce a silhouette image. The second possibility is to threshold the images prior to averaging, and then take the median value of the series of binary values. This latter method has the potential advantage of requiring the manipulation of less data, which, in certain situations might be faster and more efficient. Both techniques were examined for this work.

The two data sets examined were the real FLIR drone imagery, and a synthetic data base generated by adding gaussian noise to silhouettes of the drone model. The initial averaging tests assumed perfect image to image registration. Specifically, the drone images were averaged

directly, without attempting to spatially center the images. The model data base similarly started with the clean silhouette in the same position in each frame, and then the noise was superimposed onto it. Both these correspond to the assumption of a fixed camera view of a stationary object; even though the image is corrupted by noise, its current position *is* recognized as being correct *on the average*.

The alternate strategy assumes that the object's position within the image cannot be specified with absolute certainty. This would be the case in tracking a moving object. In this situation, to accomplish greylevel averaging it is necessary to reregister the object within the image, so that all objects are spatially aligned. The normalization process for moments eliminates translational and rotational discrepancies between images, and thus the moment averaging results perform this *reregistration* in the moment domain for all results.

For the greylevel averaging, the reregistration was assumed to require only translational correction and not rotational alteration. This would be realistic in situations where the multi-image sequence could be obtained at a rate substantially greater than the rate at which the observed object changes course. The technique was to select the first image in the sequence, and then correlate its silhouette with the silhouette of all the other images in the data base, correcting the object position in these latter images to agree as closely as possible with the first image. The greylevel averaging was then performed.

For the model data, the problem of misregistration by less than one pixel was directly introduced. Prior to the addition of noise, four ideal images were generated that varied in registration from the original model by, respectively, zero pixels (the original), one half pixel to the right, one half pixel downward, and one half pixel to the right and downward. These *subpixel* shifts were accomplished by performing full pixel shifts in a double sized image, and then shrinking the image by polling the contents of a two by two window. The result was said to be object if more than two pixels in the window were object. Ten noisy versions of each of these ideal images were then created, to give a 40 image data set. These images were

then reregistered by comparison to an arbitrarily selected reference image from the data set.

3.2. Averaging Results

The results of the first experiment could be labelled *perfectly registered* data, since the camera and object were identically aligned for each image. Thus, if the image disturbances tend to vary uniformly about the true value, we would expect averaging in greylevel domain to be effective for noise reduction. The location normalization for standard moments reacts to the slight image to image shape variation due to noise, and consequently the ideal registration is lost in the moment domain. This effect puts moment averaging at a disadvantage for this data set.

The results for moment and greylevel averaging without and with aspect ratio normalization are shown in figures 3.2 and 3.3, respectively. The greylevel averaging was performed using the actual greylevel for the real drone data, and using silhouettes for the model data. The process of aspect ratio normalization can be viewed as a distortion of the object's dimensions to produce equal *moments of inertia* about the 2 coordinate axes. The original aspect ratio is retained as a feature vector element.

The reregistered results are shown in figures 3.4 through 3.7. The figures 3.4 and 3.5 give the results for the averaging of greylevel images, and the remaining two figures show the results for the averaging of silhouette images. Again both aspect ratio and non-aspect ratio normalized results are presented.

3.3. Summary of Image Averaging

For data which is perfectly registered, our results indicate that at best moment domain averaging is comparable to greylevel averaging. This is to be expected, since the necessary discretization of an object's relocation in image space is not a disadvantage. However, since an object can be continuously relocated in the moment domain, it would be expected that imagery with fractional pixel displacements might be more successfully averaged in moment space. A very significant consideration is that object relocation in any realistic situation is

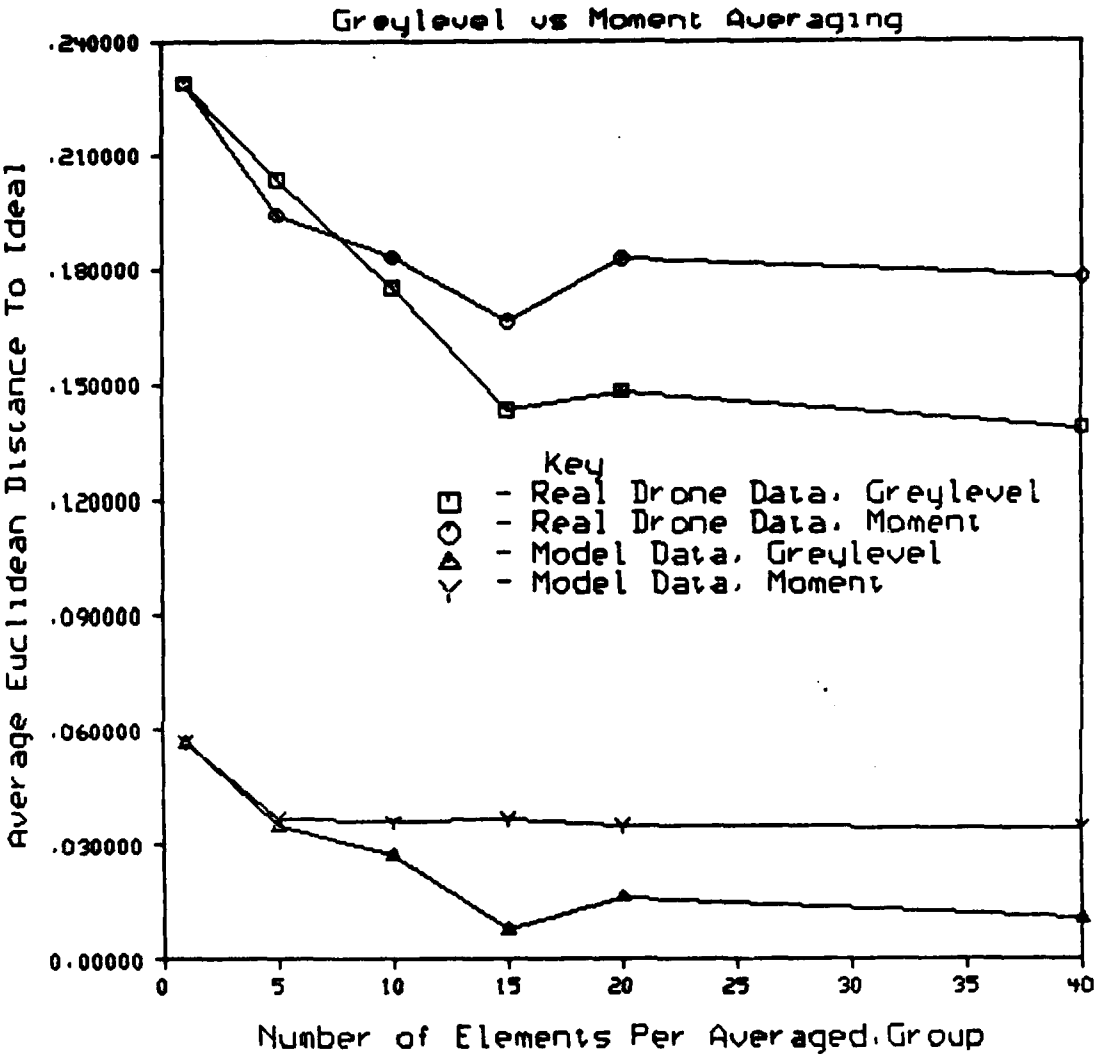

Figure 3.2.
Perfectly registered data results for moment and greylevel averaging of standard moments.

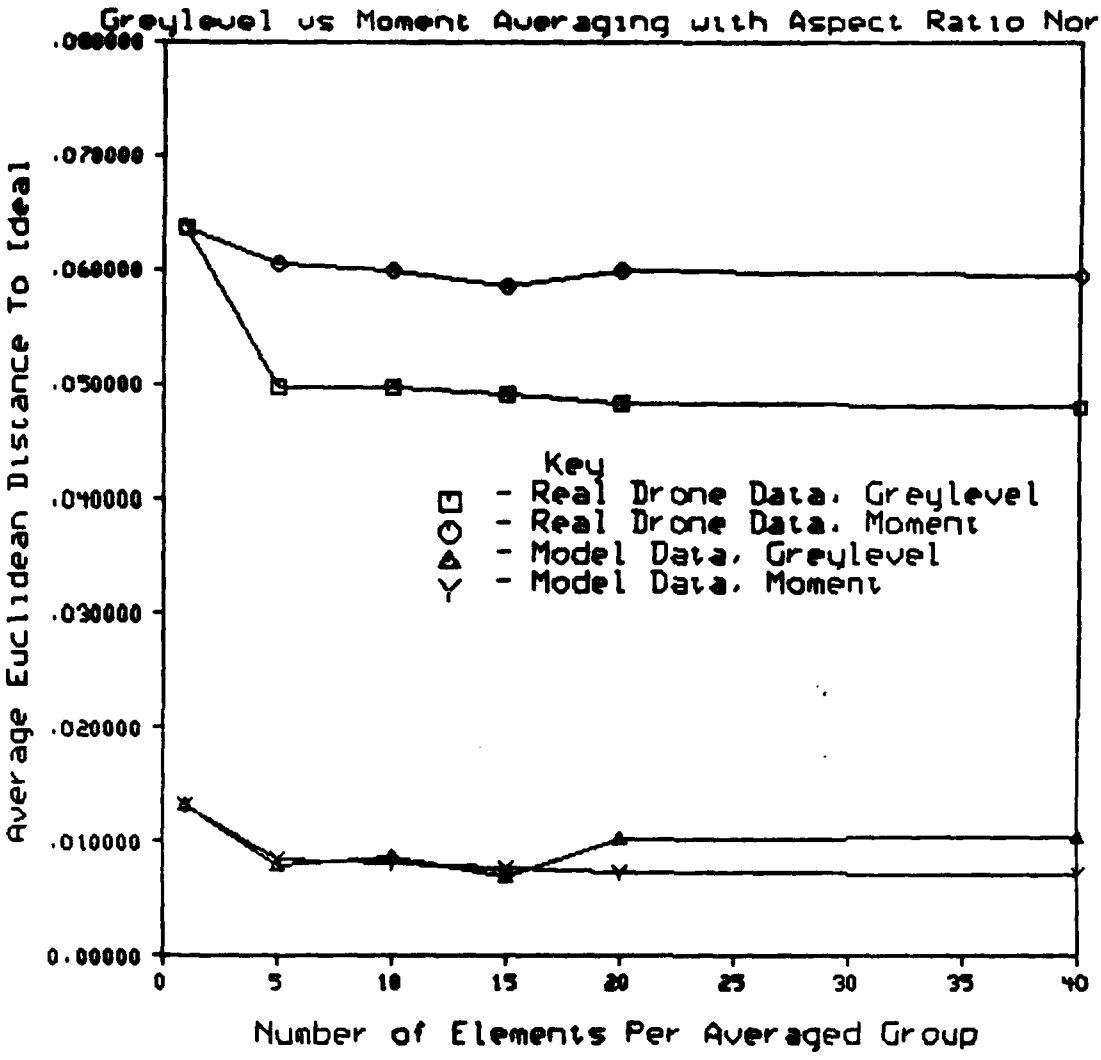

Figure 3.3.
Perfectly registered data results for moment and greylevel averaging of aspect ratio normalized moments.

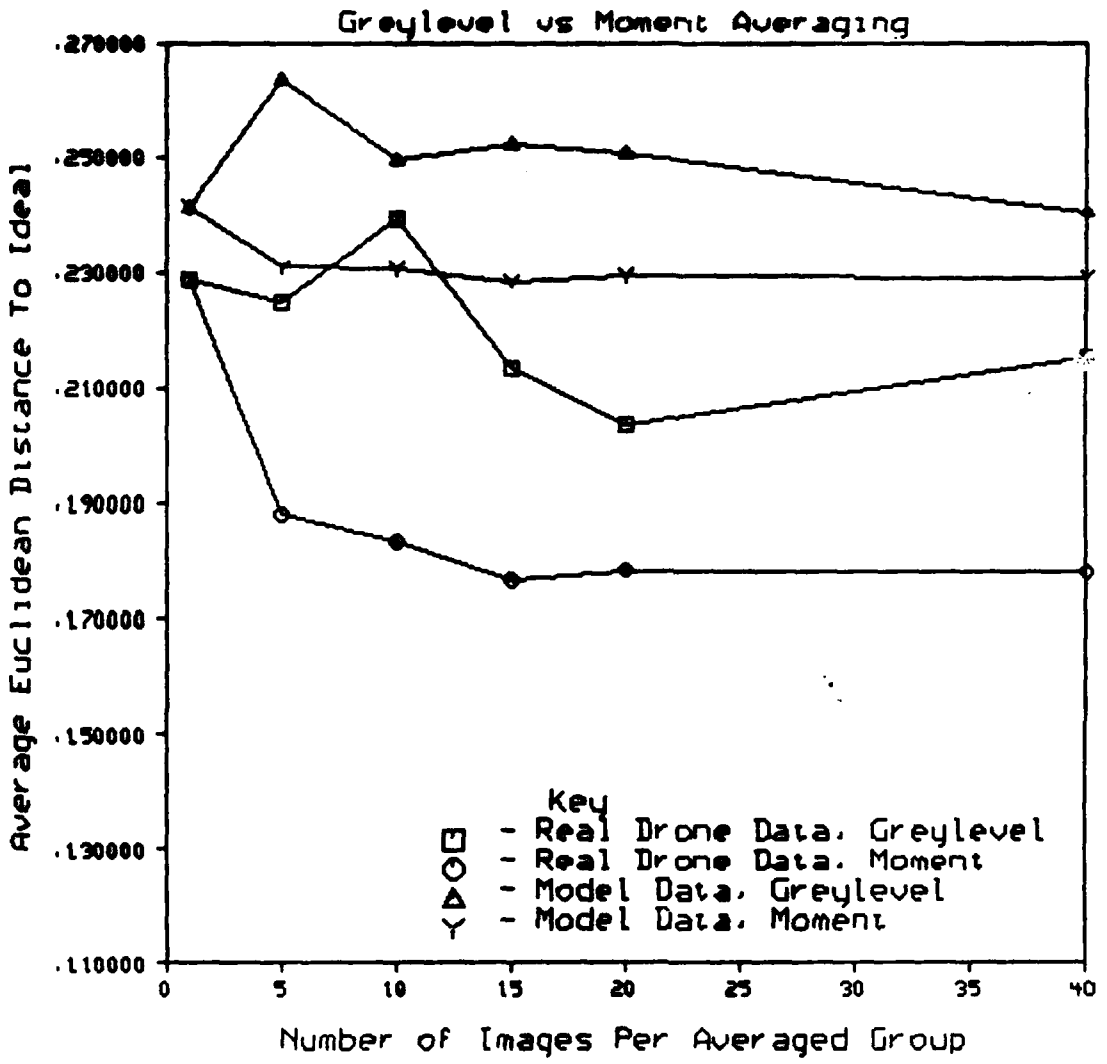

Figure 3.4.
Results for moment and greylevel averaging of the standard moments of the reregistered data. The greylevel averaging was performed using actual greylevel images.

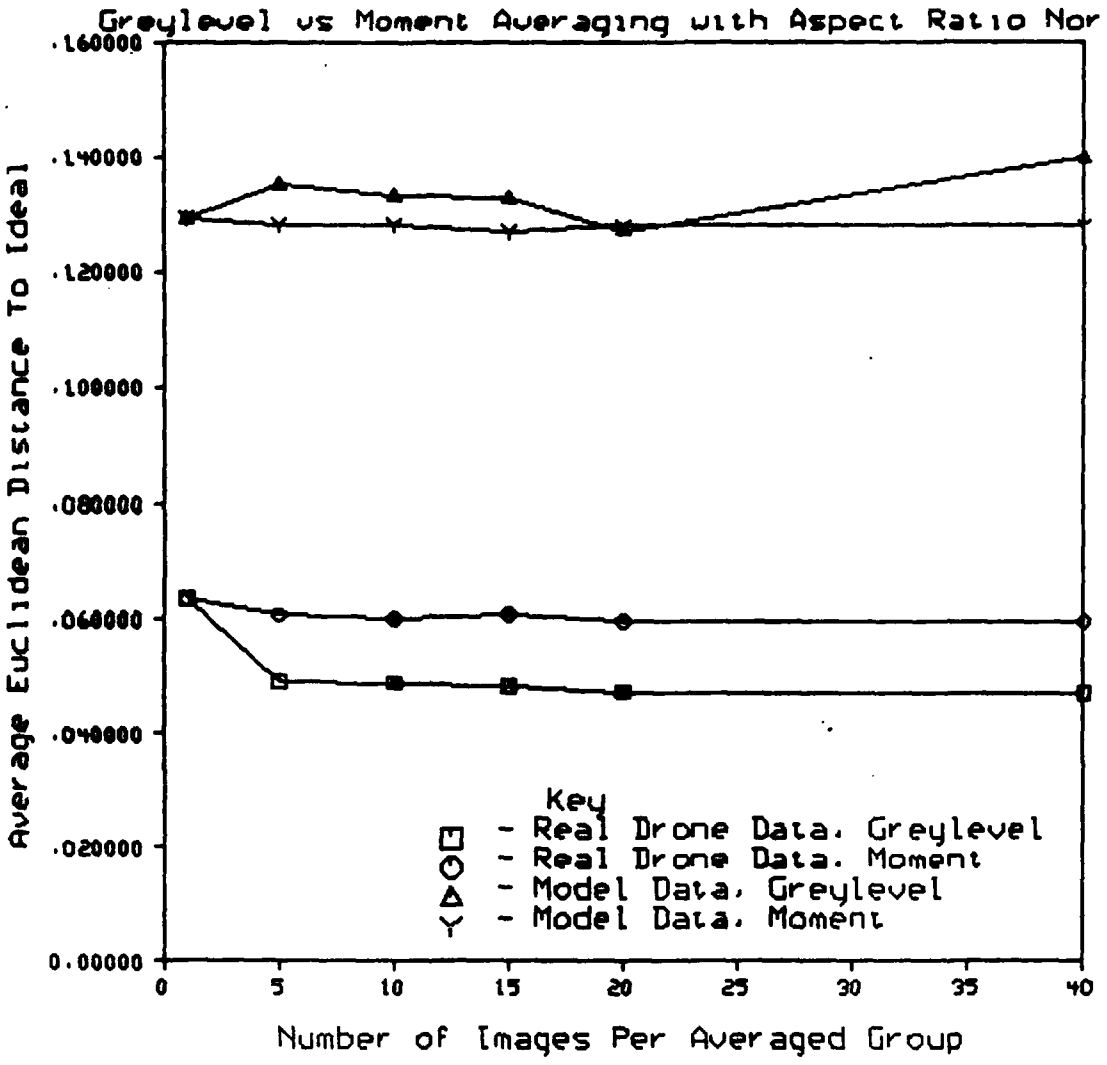

Figure 3.5.
Results for moment and greylevel averaging of the aspect ratio normalized moments of the reregistered data. The greylevel averaging was performed using actual greylevel images.

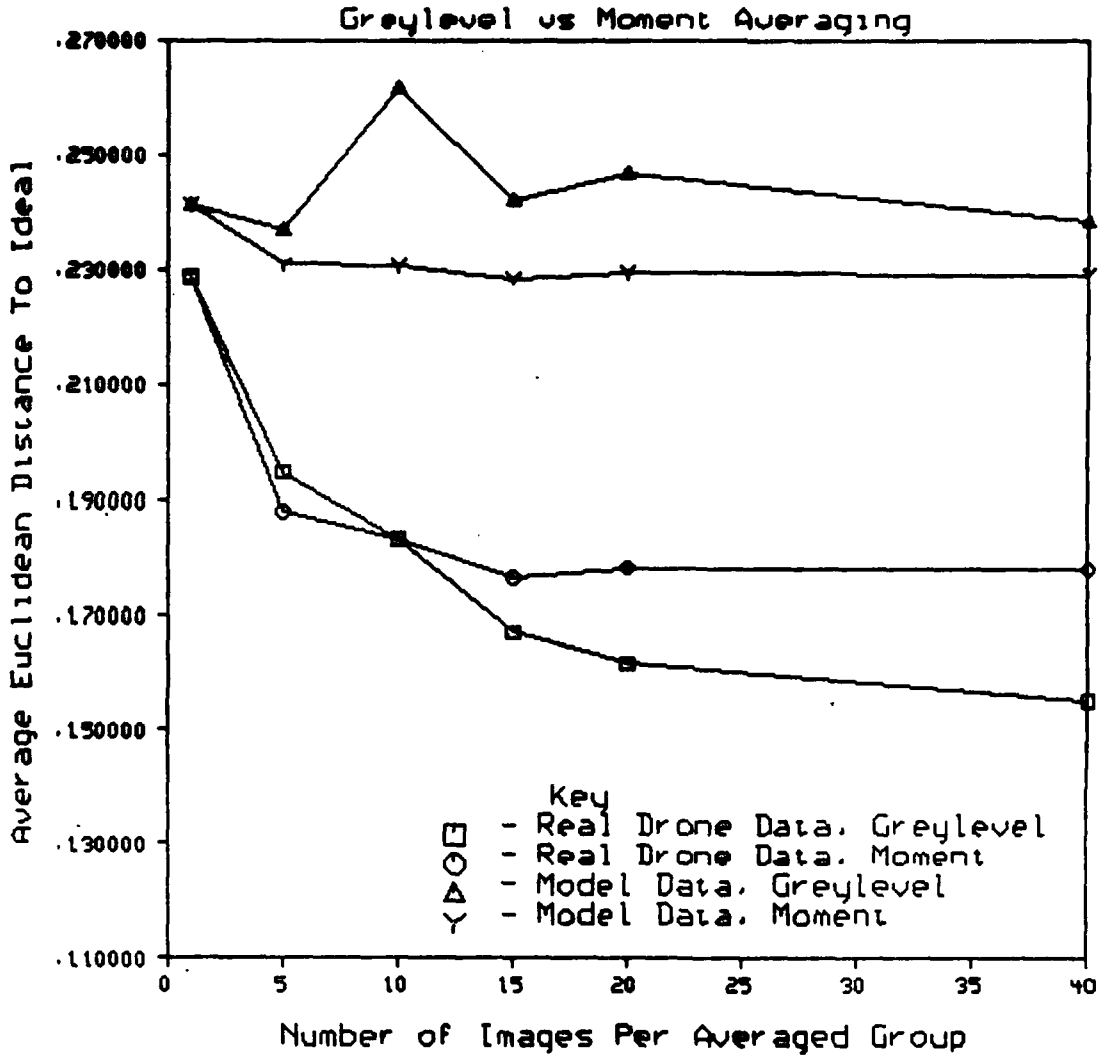

Figure 3.6.
Results for moment and greylevel averaging of the standard moments of the reregistered data. The greylevel averaging was performed using binary silhouette images.

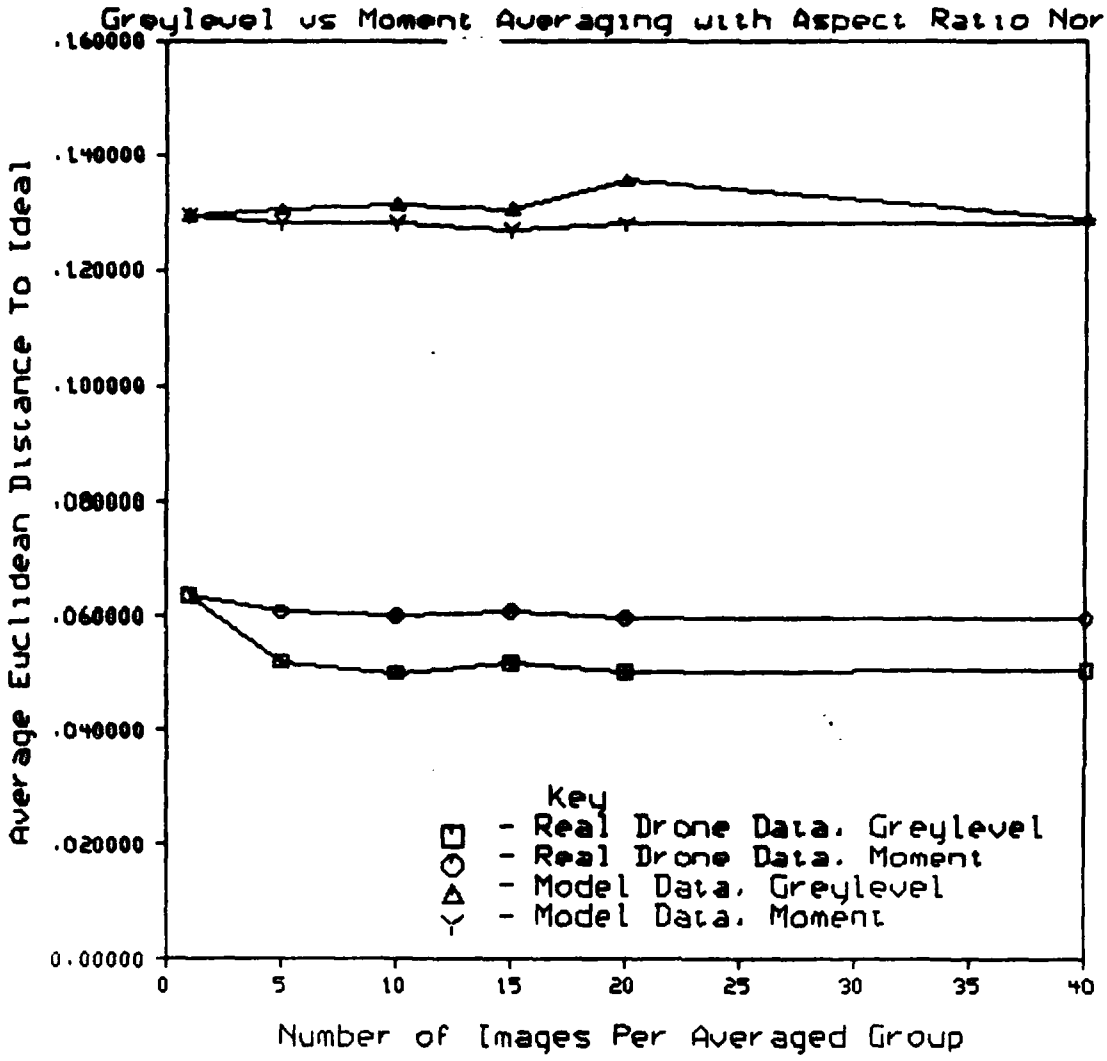

Figure 3.7.
Results for moment and greylevel averaging of the aspect ratio normalized moments of the reregistered data. The greylevel averaging was performed using binary silhouette images.

much less computationally costly in the moment domain than in the image domain.

The results for the reregistered data show that moment averaging is, in a number of cases, much more successful than greylevel averaging. In particular, for averaging actual greylevel images, without aspect ratio normalization, moment averaging produces superior results. The introduction of aspect ratio normalization in all cases produces better results for greylevel averaging of the real drone data.

The effects of aspect ratio normalization on these results is noteworthy. For the real drone data, any improvement resulting from moment averaging seems to be obviated by the process of aspect ratio normalization. This does not hold true for the model data. In general, the overall improvement in data quality from averaging in either moment or image domain is not as notable with aspect ratio normalization as without.

Since earlier results have indicated that superior object identification results are obtained from aspect ratio normalized moment sets (see also the next section of this report), the conclusion is that aspect ratio normalization reduces the sensitivity of the moments to noise. The mechanism that achieves this is quite basic; the further a noise pixel is located from the centroid of an object, the greater its relative effect on that object's moments. Aspect ratio normalization contracts protruding portions, and expands compressed regions, of an object that make the aspect ratio non-circular. Thus, for the drone image, noise on highly corruptable image areas, such as the wing tips, may have a much smaller relative effect on the moments after aspect ratio normalization.

A simple experiment was performed to verify this hypothesis. A comparison by Euclidean distance was made between the ideal model's moments, and the moments of the model with one pixel removed on a wing tip, and with one pixel removed from an extreme point of the body. The results showed that the Euclidean distance to the model missing the wing pixel was four times larger than the distance to the model missing the body pixel, for standard moments. But these results were reversed for aspect ratio normalized moment sets; the ratio was now 1 to 3, instead of 4 to 1. Hence, after aspect ratio normalization, the wing

tip pixel error produced a much smaller relative change in the moment values.

The effect of group size on averaging improvement seems fairly uniform for both greylevel and moment averaging. In most cases, an averaging of 5 frames seems to produce most of the improvement. In all cases very little, if any improvement is realized by averaging more than 15 images.

A tertiary conclusion is that averaging of silhouette images is more effective for noise reduction than averaging the unthresholded greylevel images. This is probably an artifact of the preprocessing that is performed on the silhouette images. This effect is particularly notable for the real drone data.

Overall, then, the results for moment averaging are better than those for greylevel averaging when there is uncertainty about the object's position within the image. Under some conditions, determined by the conformation of the object, aspect ratio normalization may decrease, or eliminate the noise reducing effects of multiple image averaging. In cases where image averaging does improve the data, an average of 5 to 15 images should be sufficient. Additionally, averaging of silhouette images for greylevel averaging produces the best results.

4. Classification Using Moment Subsets

The primary purpose of a feature vector representation of an image *is* to achieve identification. The greatly reduced dimensionality of the feature vector, as compared to the original image, makes possible techniques such as library matching for object classification. Clearly, in the spirit of this approach, it is desirable to use the fewest descriptive numbers possible to accurately characterize an image.

This section considers the possibility of image classification using various complete moment set (*CMS*) subsets. This examination is intended to give increased insight into the value of moments with certain common characteristics for uniquely identifying samples from a six airplane data set.

4.1. Experimental Methods

The classification experiments carried out here examined moments that were normalized for size, translation, and orientation. A distinction was also made for normalized moments with, and without, aspect ratio normalization. The experimental technique for classification used a nearest neighbor Euclidean distance metric, where all redundant or trivial information was removed from the feature vectors. In the case of aspect ratio normalized moment subsets, the aspect ratio was included in all subsets for classification purposes. All images contained only silhouette data.

Moments up to order 6 were considered in these trials. A set of 252 *unknown* views of each of 6 airplanes were classified using a *known*, or library, data set of 500 views of each of the same 6 airplanes. The classification was deemed correct if the unknown view was labelled as being the appropriate airplane.

The design of the classification task was similar to that described in section 2, except that a more regimented procedure was defined for determining the viewing angles. The technique involved approximating an even partitioning of the surface of a sphere by using views on the faces of a polyhedron. These views were used as the known library views, and then a set of

worst case unknown views were generated using viewing positions located at the polyhedron's vertices. The views are termed worst case, since they represent the interstitial point furthest from the 5 or 6 near neighbor library views.

Further normalization of the moment sets was achieved using the technique of *variance balancing*. This involved determining the standard deviation of a given moment value over the entire range of library views. All moment values in both the library views and the unknown views were then normalized by their respective standard deviation. This step was performed after all other moment set normalizations, and served to deemphasize highly variable moment values in the Euclidean distance measurement.

4.2. Clean Subset Results

The main analysis was performed on noiseless, or *clean* data. The most obvious moment subsets to consider are the CMS subsets of our order 6 CMS. Figure 4.1 shows these results for CMS of order 3, 4, 5, and 6, with, and without, aspect ratio normalization. One observation is that the change from an order 4 to order 5 CMS gives an improvement in performance that is disproportionate to the increase in feature vector size. Conversely, there is little or no improvement in classification success for an order 6 CMS, as compared to an order 5 CMS. In all cases, aspect ratio normalization gives superior results to those for moment sets without this information. This result will be reenforced by the other results presented.

The next class of subsets tested were those composed solely of moments of a selected order. Orders 3, 4, 5, and 6 were considered individually; the results are presented in figure 4.2. Aspect ratio normalized moments again show superior results to moment sets without this treatment. Most importantly, however, it should be observed that the odd order moments give better results than the even order moments for this small sample. The order 5 moment subset gave results equivalent to the order 3 CMS classification results, where both are of similar size.

The results in figure 4.2 help explain the CMS results. The conclusion would seem to be that, for at least this silhouette data set, the order 5 moments are very important to the

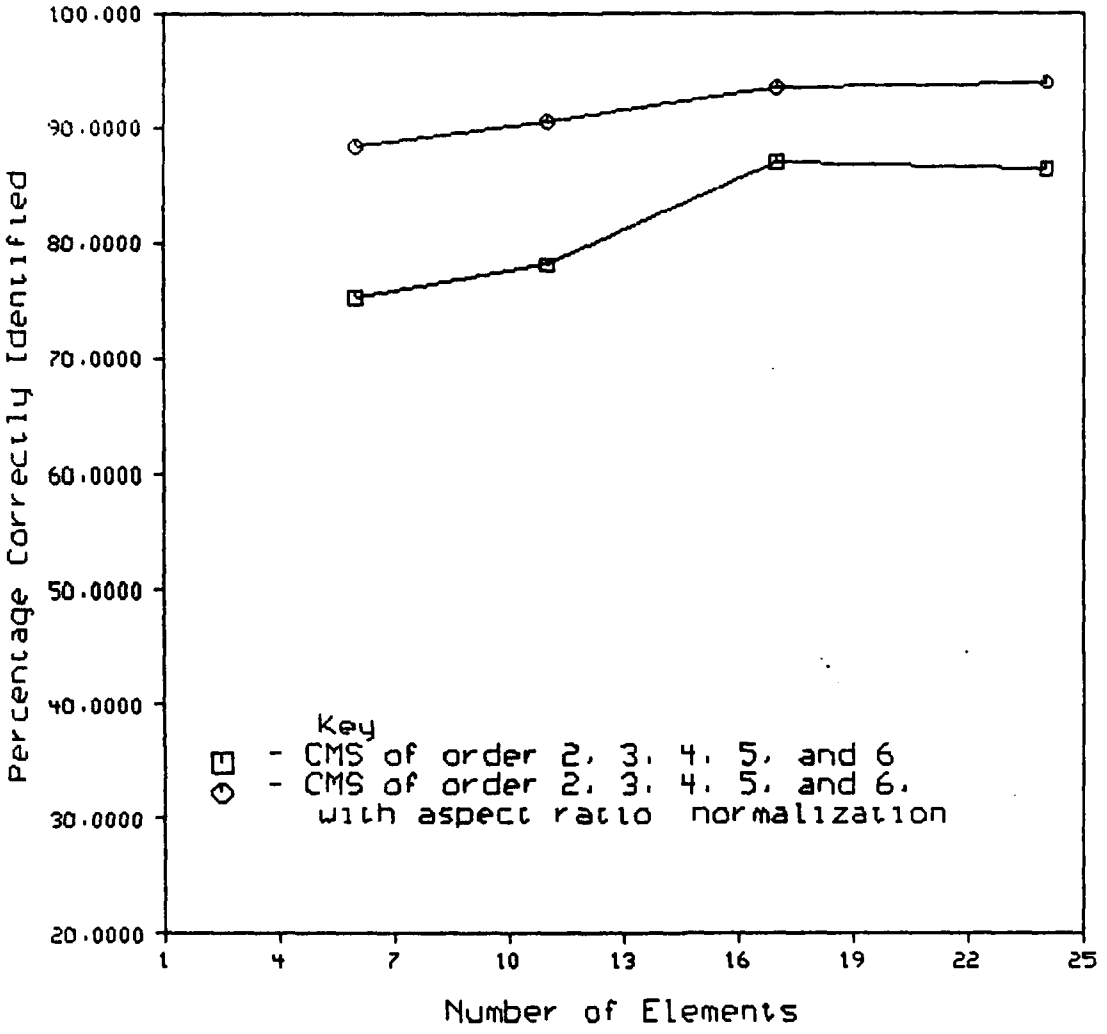

Figure 4.1.
Classification results for complete moment sets of order 3, 4, 5, and 6. Both standard moment and aspect ratio normalized moment results are shown.

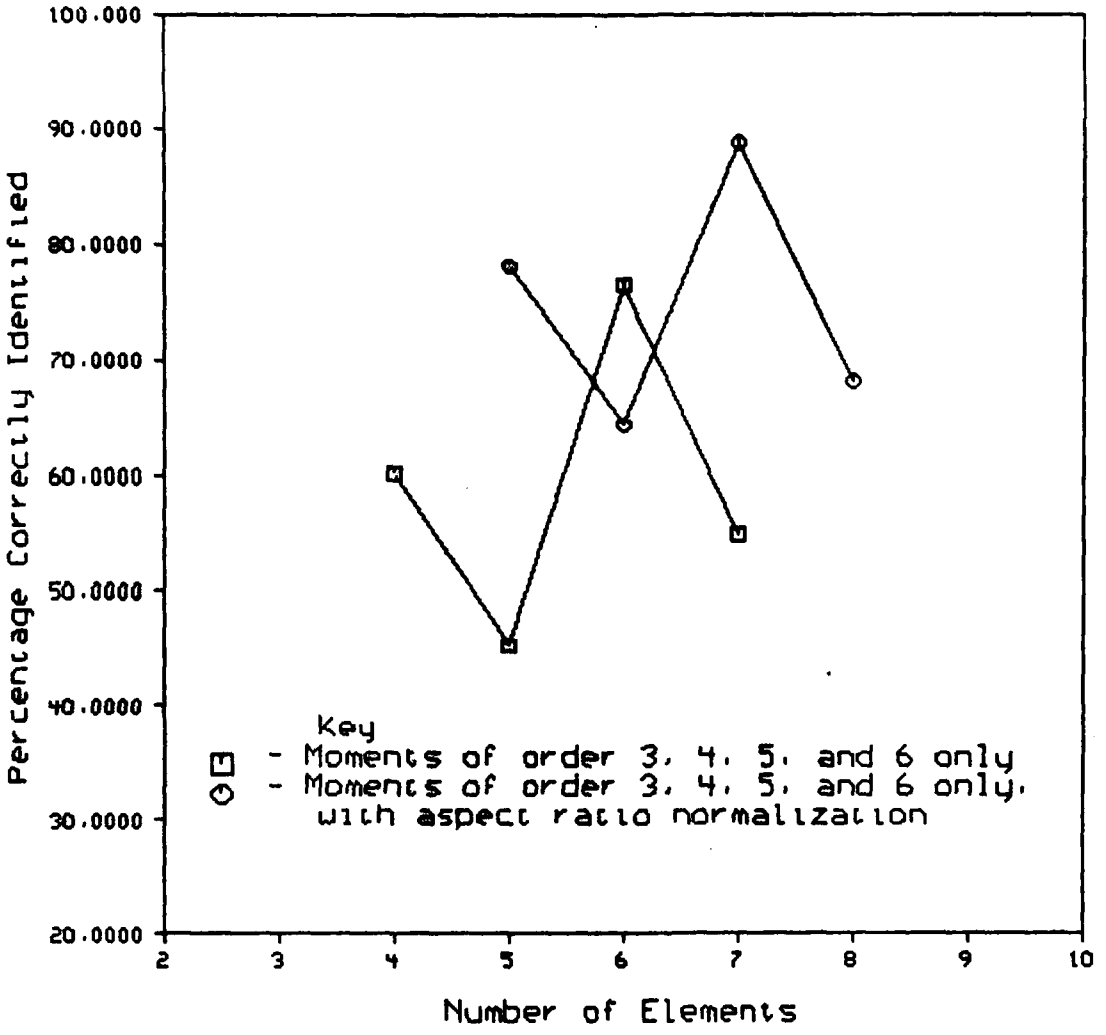

Figure 4.2.
Classification results for single moment orders. Results are shown for orders 3, 4, 5, and 6 for standard and aspect ratio normalized moments.

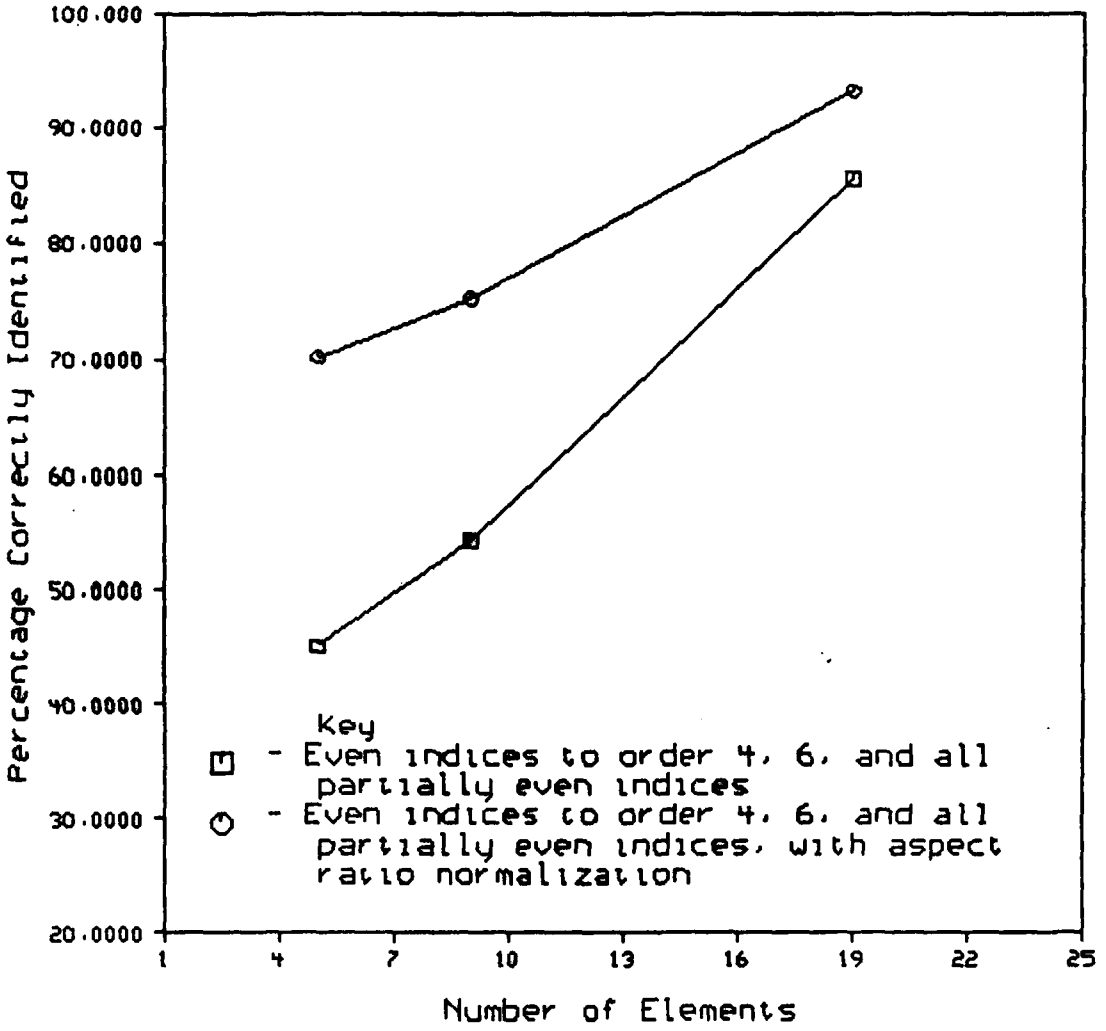

Figure 4.3.
Classification results for moments with even, and partially even indices. Both standard moment and aspect ratio normalized moment results are shown.

classification success rate. Hence, a large improvement in classification success is observed when an order 5 CMS is compared to an order 4 CMS. A possible generalization, which would require more experimentation to validate, is that odd order moments are more important to classifying silhouette data than even order moments. In general, then, it is possible that adding the next odd order of moments to a given feature vector could be the most efficient way to increase classification success.

Figure 4.3 shows the results for moments with even and partially even order indices. Due to the nature of the moment calculations, the moment values with both indices even numbers tend to have the largest values for a given order. Thus it is interesting to see if this relative size difference has any bearing on the classification success rate.

Figure 4.3 indicates that strictly even indexed moments have no notable success for classification. Adding in moments with only partially even indices dramatically improves the results, which now approach the success level of a similarly sized CMS, though still not quite as good. These results suggest that moments with even indices do not have any special properties that single them out as desirable for classification purposes.

The results for moments with strictly odd indices are very similar to those for strictly even indices. Table 4.1 shows these and some other results. The expression *axis projection* refers to all moments that have one specific index equal to zero. This is tantamount to determining the moments of a projection of the image onto one of the coordinate axes. After rotation normalization, the long axis of the object silhouette is oriented along the x axis. Table 4.1 indicates that the moments of this long axis projection are much better than the y axis projection for classification purposes. Aspect ratio normalization eliminates this discrepancy between the projections, and adds a new piece of information in the form of the aspect ratio, and consequently reduces the gap in success between the two classification efforts. Classification using both axis projection moments shows improvement that at best is of the order that would be expected from doubling the feature vector size.

Table 4.1

Moment Subset Classification Results*		
Subset	Number of Elements	Results
odd moment indices	5 (6)	44.1 (50.0)
x axis projection	5 (6)	47.4 (73.2)
y axis projection	5 (6)	36.4 (70.1)
x and y projections	10 (10)	65.5 (81.9)

*aspect ratio normalized results in parenthesis

In general, the subsets summarized in table 4.1 do not produce exceptionally good classification results, the results being clearly inferior to the CMS results, and equivalent or slightly better than the strictly even moment indices tests. Again, the results reassert that aspect ratio normalization is, in all cases, worthwhile.

The experiments above reenforce each other. For example, the poor results for both all odd and all even moment indices make the poor results for even order moments not at all surprising, since *all* moments in an even order will have both indices odd or both even. Similarly, an odd moment order will have all mixed indices. Thus the tests of strictly even, and even and partially even indexed moments in figure 4.3 simply means we are comparing a classification of even order moments to one of even *and* odd order moments. The only reason to compare classification success for moments with both indices even, and both odd, is to attempt to isolate why even order moments perform so inadequately.

The final tests involve actually applying the results of the above experiments. Table 4.2 shows some comparisons of different moment subset classifications. The entries are ordered by feature vector length, and all are aspect ratio normalized. The intent is to attempt classifications using just odd order moments, with various pieces of additional information also included. The results for several CMS sizes are shown for comparative purposes.

Table 4.2

Aspect Ratio Normalized Moment Subset Classification Results		
Subset	Number of Elements	Results
CMS of order 5	17	93.5
Aspect ratio, orders 2, 3, 5	12	93.3
Aspect ratio, orders 3, 5	11	90.5
CMS of order 4	11	90.5
Orders 3 and 5	10	90.1
CMS of order 3	6	88.4

Table 4.2 points out a previously unvoiced assumption: the order 2 moments, though even, are very important for classification success. The reasoning is that the low order moments contain very fundamental information about an object, such as area, location of the centriod, and approximating ellipse. Higher order moments would then contain more detailed information, what might heuristically considered to be *high* frequency variation. As an extreme example, the moments of a circle and an ellipse, after full normalization (including aspect ratio), might differ only by their aspect ratios, which are a function of their order 2 moments. In table 4.2, adding the order 2 moments to a feature vector already containing the aspect ratio and the order 3 and 5 moments dramatically increases the classification success rate, at the cost of only 1 element more in length.

4.2.1. Noisy Subset Results

An examination of the same subsets as described above was performed for noisy image data. The noise consisted of a boundary perturbation of the silhouette image by randomly removing some of the object's edge pixels. The image resolution was also reduced by 25 percent. Figure 4.4 shows some sample clean and noisy silhouettes. This type of noise is not as harsh as some types examined, but is a reasonable approximation to the type of noise a range finding imaging device might produce.

In general, the previous results are directly supported. The noisy trials were only performed for aspect ratio normalized feature vectors, since all experimentation to date has indicated this type of moment normalization improves classification success, and hence is of the most interest to us here. The classification results for the noisy data seem to fall universally about 2 to 5 percent short of the clean results, and consequently support the clean result trends.

4.3. Summary

As the classification success becomes very high for large moment based feature vectors, it becomes increasingly harder to improve the success rate without adding comparatively large

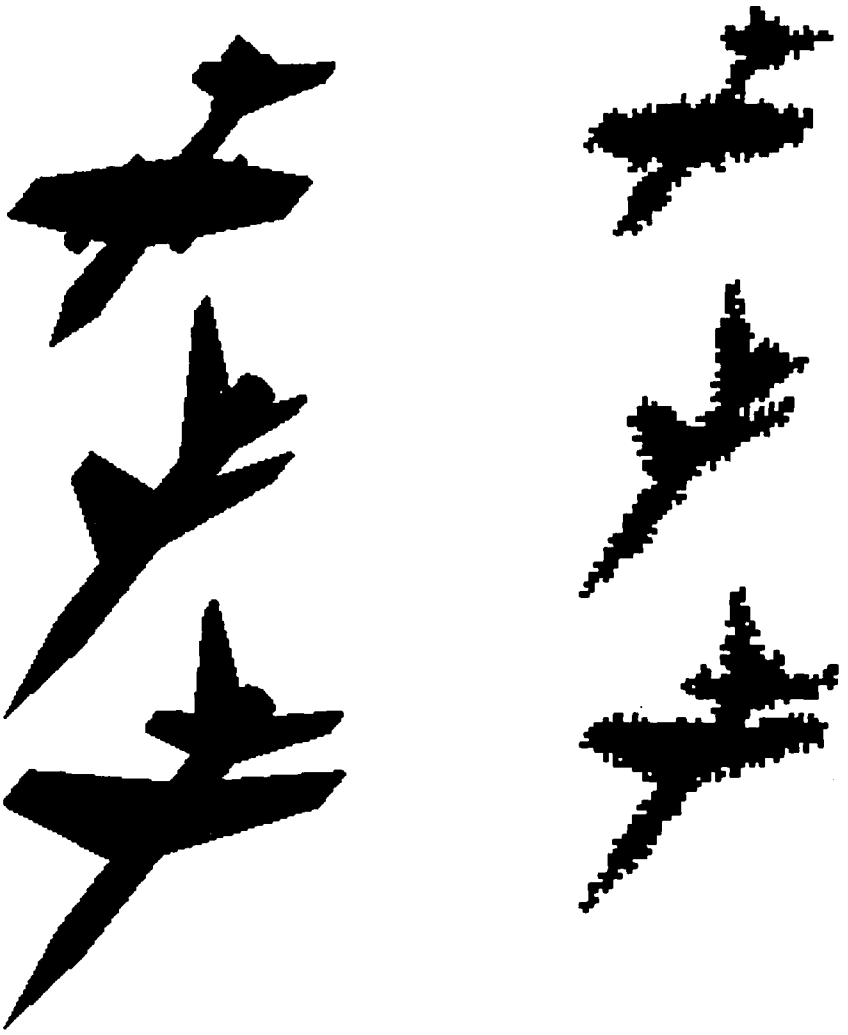

Figure 4.4.
Sample clean and noisy test images. Clean versions of three of the six airplanes are shown on the left, and the reduced resolution noisy versions of the same views are shown on the right.

numbers of elements to the feature vectors. The intent of this research was to determine if certain characteristics could be identified that were common to moments that had either a positive or negative correlation with classification success. The logical supposition is that, at a high level of classification success, some moments may be adding ambiguity to the classification attempts, rather than clarity. Consequently these moment values should be eliminated from the sample.

On the most superficial level, the results here suggest that a CMS is probably the best choice for the most successful and economical classification attempts. It also is apparent that aspect ratio normalization improves the results for any given silhouette moment set. Most of the remaining moment subsets examined produce similar results that are inferior to those for the CMS.

A more subtle result is that, for the small sample size examined here, odd moment orders are superior to even moment orders for classification purposes. Specifically, for the aspect ratio normalized moments we find that the moments of order 5 give better results than a CMS of order 3, where both contain the same number of elements. It might be worthwhile to try and extend this result at some point by examining higher moment orders, but the initial conclusion must be that a CMS based classification effort should contain an odd order as the maximum order in the set for best results.

Similarly, if information must be eliminated from a CMS to reduce the feature vector size, it is probably best to eliminate the even moment orders, corresponding to those moments with all even or all odd indices. Trying this experiment shows that a feature vector containing aspect ratio normalized moments of orders 2, 3, and 5, and the aspect ratio, exhibits a reduction in classification success of only .6 percent from an order 6 CMS, and yet contains only one-half as many elements. Hence a 29 percent size reduction of the order 5 CMS feature vector, and a 50 percent reduction of the order 6 CMS, is achieved with a disproportionatly small decrease in classification success.

The suggestion would seem to be that, in general, not all moment values are of equal import for object recognition. In situations where the views are more typical or consistent than in the 6 airplane experiment presented here, it may well be possible to identify very specifically the moment values necessary for object identification from silhouette imagery.

5. Image Reconstruction from Moments

The main application of moments in image processing has been for object identification. For example, the reduced dimensionality of the moment based feature space makes possible the use of library matching techniques. An interesting question is how well can the original image be recreated from a set of its moments. The theory suggests that a perfect reconstruction of a continuous image requires an infinite set of moments. However, a suboptimal reconstruction may still retain the significant traits of the original.

A further consideration is that, when reconstructing silhouette images, there are only two intensity values. Hence, it is possible that such *a priori* knowledge would allow a reconstruction far more accurate than might initially be anticipated. Since the data analysis in this report is based on silhouette imagery, this type of reconstruction is of particular interest.

Teague [4] considered possible techniques for image reconstruction using an inverse moment transform. He identified two methods, one using standard moments, and one using Legendre polynomial-based moments. The first approach produces a number of sets of coupled equations. These equations can be solved to give the coefficients of a polynomial as a function of all moments of order n and lower, where n is the degree of the polynomial. This polynomial defines a continuous image function that has the same moments as the original, up to the order utilized in the reconstruction.

Key problems for this approach are the solution of coupled sets of equations, and the necessity of recomputing the polynomial coefficients for a larger order moment solution. Teague's [4] second option suffers from neither of these problems, and is examined here, along

with an iterative technique that dramatically improves silhouette reconstructiom results.

5.1. Image Reconstruction from Legendre Moments

5.1.1. Legendre Moments

The Legendre moments of an image are generated similarly to conventional moments. The difference is that the basis functions for Legendre moments are Legendre polynomials, rather than the monomials of x and y found in the standard moments. Specifically [4]:

$$L_{mn} = \frac{(2m+1)(2n+1)}{4} \int\int f(x,y)\, P_m(x)\, P_n(y)\, dx\, dy \qquad (5.1)$$

where L_{mn} denotes the Legendre moment of order $m+n$, and $P_m(x)$ refers to the Legendre polynomial in x of order m. The Legendre moments can also be generated as linear combinations of the conventional moments [4]. Furthermore, a complete moment set of the Legendre moments contains the same total amount of information as a conventional moment set of the same order.

However, Legendre and conventional moments are not entirely equivalent. In particular, Legendre polynomials are orthogonal in the interval [-1,1], which means that the moments of an object can be determined using orthogonal basis functions, if the object is scaled to fit dimensionally into these limits. The result of using orthogonal Legendre moments in an image reconstruction attempt is that the algebraic equations to determine the inverse transform polynomial coefficients are no longer coupled. Additionally, these polynomial coefficients no longer change when the moment order considered changes, allowing easy precomputation of the necessary values.

The reconstructions resulting from Legendre moments and the conventional moments are identical, since a complete moment set of either to order n contains exactly the same information. The Legendre moments simply present the information in a particularly tractable

and efficient format.

5.1.2. The Reconstruction Technique

The polynomial defining the image reconstruction for a complete, but finite Legendre moment set can be expressed as [4]:

$$g(x,y) = \sum_{N=0}^{N_{max}} \sum_{n=0}^{N} L_{N-n,n}\ P_{N-n}(x)\ P_n(y) \qquad (5.2)$$

The Legendre polynomials in the above are normalized for the [-1,1] interval [4]. Equation (5.2) yields a continuous function, $g(x,y)$, with the same moments to order N as the original discretely sampled image function. If $g(x,y)$ is discretely sampled, and additionally thresholded into a silhouette image, the moments will no longer match those of the original.

The problem of perfectly reconstructing a binary valued, discretely sampled image directly from this method is impossible, since the original image violates necessary continuity assumptions. Consequently, even increasing the order of the moment set used in the reconstruction will not guarantee a good result. In an effort to compensate for this problem, an iterative scheme using error feedback was devised to help reconstruct silhouette images.

The fundamental approach in this iterative scheme was based on the fact that the moment transform is a linear operation. This means, as an example, that the moments of the difference between two images are the same as the difference of the two image's moments. Using this information, an error image can be reconstructed from the moment set error, and then subtracted from the current reconstruction to enhance its accuracy.

The actual algorithm, then, produces an iterative image reconstruction by negative feedback of the error. An outline of this procedure is:

(1) Perform the basic image reconstruction from the Legendre moments to produce *image(0)*

(2) Set $i = 1$

(3) Produce a thresholded image, *binaryimage(i-1)*, by calculating the threshold point that yields the same image size as in the original image (specifically equal L_{00})

(4) Calculate the Legendre moments of *binaryimage(i-1)*, calling these moments *LM(i)*

(5) Determine the difference between the moments of the original image and *LM(i)*, and reconstruct an image, *diffimage(i)*, from this difference

(6) Calculate *image(i)* = *image(i)* + β * *diffimage(i)* where β is a feedback coefficient between 0 and 1

(7) Increment *i* for the next iteration, and go to step 3

By choosing an appropriate value for β, it is possible to gradually converge towards an improved silhouette image reconstruction.

5.2. Results

A series of silhouette images were reconstructed from moment sets of various orders. Figures 5.1 to 5.3 show reconstruction attempts for a top view of the BQM-34 Drone model, using complete moment sets of order 4, 6, and 12. These represent reconstructions using 15, 21, and 91 numbers, respectively. The image size used is 64 by 64, so this can be viewed as an attempt to recreate the original image using from a 273 to 1 to a 45 to 1 compression of the original data set. Figure 5.4 shows a graph of the pixel based error reduction accorded by the iterative technique applied to the BQM-34 image.

Figures 5.1 through 5.3, and 5.5 through 5.9, show the results of the iterative process on image reconstruction. The upper left corner image in all these figures is the original silhouette image. Directly below that is the reconstruction using Teague's method directly, and then the lower left image is the result after 250 iterations of the reconstruction algorithm. The right side of these figures show the result of successive iterations, with Teague's reconstruction first, and then the first 31 iterations following. The progression is from top to bottom, and then left to right. These figures are intended to give both a visual illustration of the change the itera-

tions produce in the original reconstruction attempt, and also an appreciation of the rate of change with different feedbacks and moment set orders.

Examination of figure 5.4 shows that, even though the result for 250 iterations is a more asthetically appealing approximation of the original, for moment sets of order 4 and 6 the absolute pixel error is not improved. Dramatic improvement is evident for an order 12 reconstruction. As a comparison case, figure 5.10 shows the same graph for the MiG aircraft. The MiG represents a fundamentally simpler geometric shape than the drone, and the reconstruction statistics reflect this fact, with both order 6 and 12 producing improved results. An order 4 moment set still shows increasing pixel error for our reconstruction attempts.

Figures 5.5 through 5.8 show reconstructions of two simple geometric shapes . Figures 5.5, 5.6, and 5.7 show, respectively, the results for reconstruction of a small rectangle using order 4, 6, and 12 moment sets. It is significant that for this simple shape even an order 4 moment set gives an excellent reconstruction after a few iterations. The iterative procedure can also use a comparitively high level of feedback, without loss of stability. Figure 5.8 shows an order 12 reconstruction of a hollow rectangular shape. For this case very low feedback was used to maintain stability of the iterations, but the reconstructed result is very accurate, despite a slight increase in complexity over the basic rectangle.

Figure 5.9 shows an order 12 reconstruction of a MiG-type aircraft. The reduced angularity of the MiG, as compared to the drone, produces a visually more representative reconstruction at both low and high numbers of iterations. Figure 5.10 is a graph of absolute pixel error for the MiG, based on order 4, 6, and 12 reconstructions. Unlike the drone, the MiG shows improvement for an order 6 moment set reconstruction, as well as for order 12. The order 4 moment set reconstruction still produces increased error for iterative attempts at improvement.

The iterations can become unstable with too large a value of β (feedback). Consequently, it was necessary to select a value for the feedback that would encourage quick convergence, while remaining stable. For the cases given here, the simple rectangle used

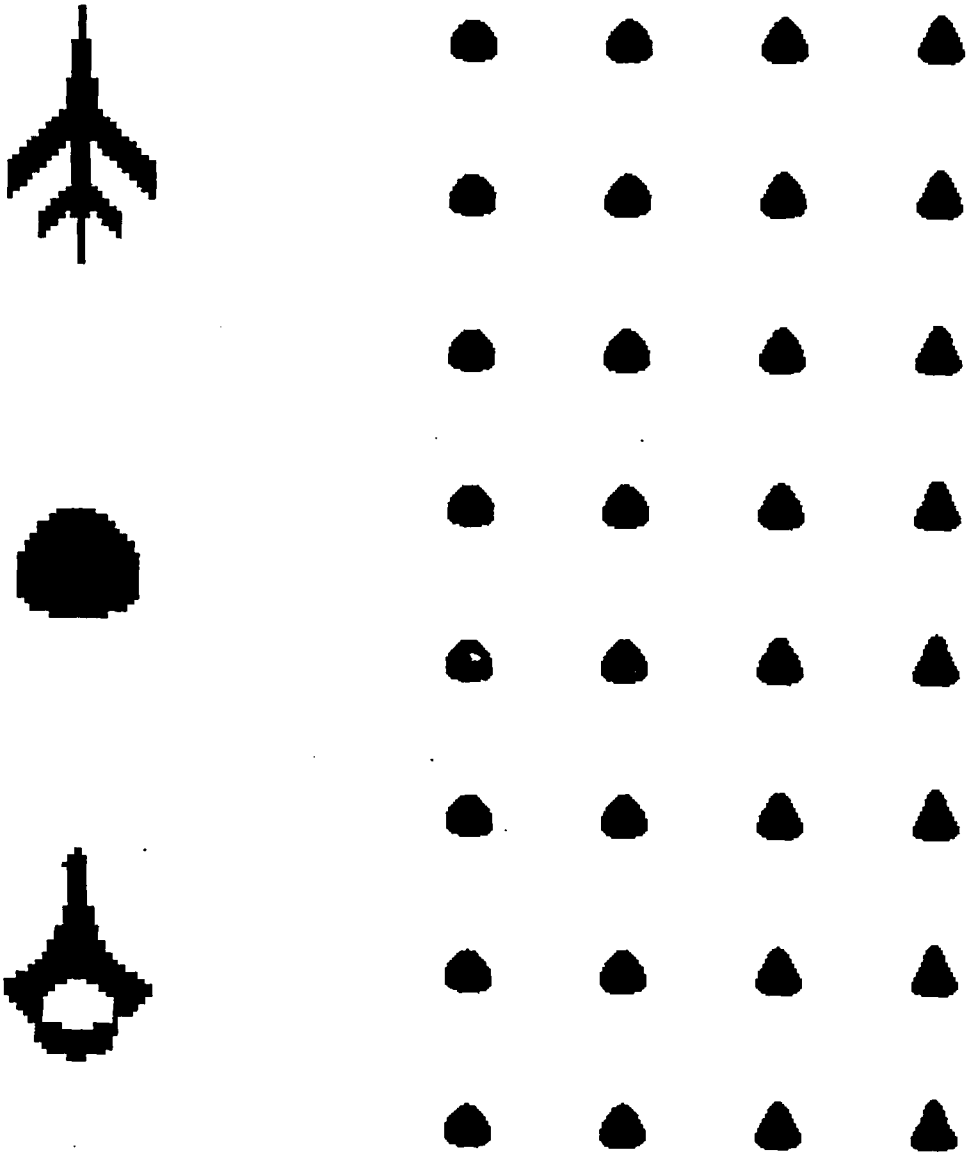

Figure 5.1.
An order 12 moment set reconstruction of the top view of the BQM-34 model, at a resolution of 64x64 pixels. The result after 250 iterations has started to show some of the features of the original.

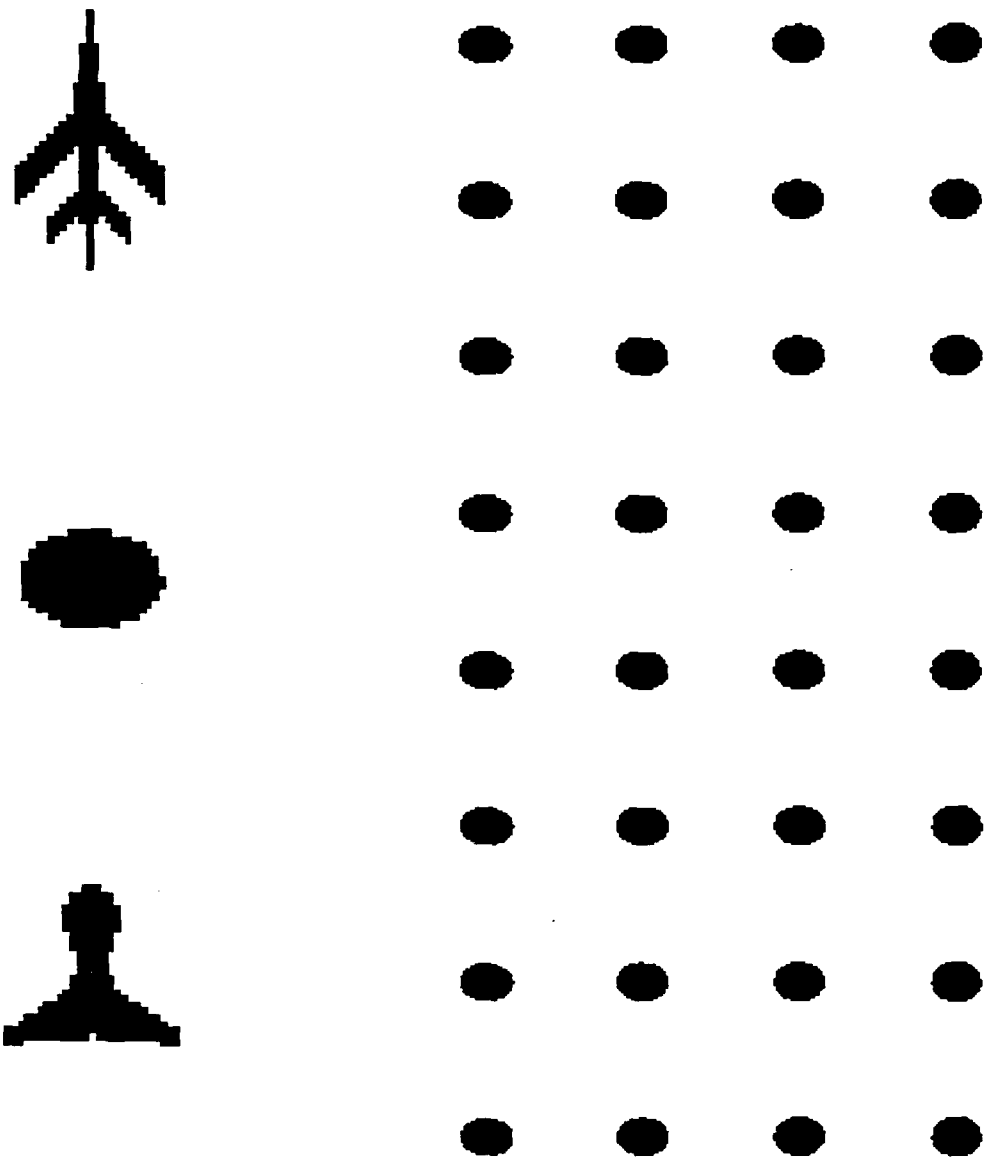

Figure 5.2.
An order 6 moment set reconstruction of the BQM-34 drone. The final result after 250 iterations appears more representative of the original than the first attempt, but in reality the pixel error rate is higher. Note the slow rate of change of the iterative reconstruction attempts with a feedback of only 2% to maintain stability.

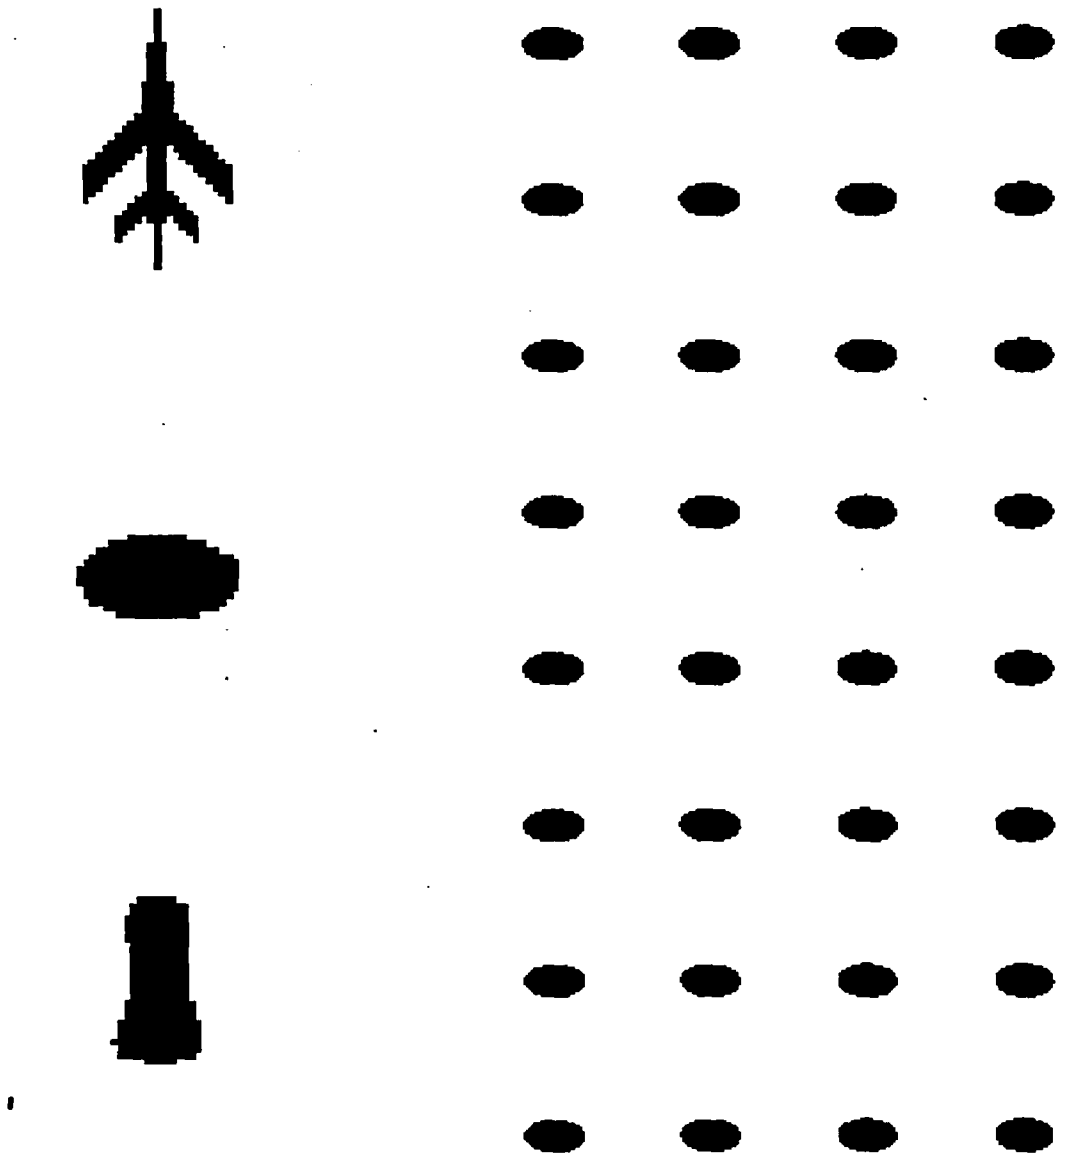

Figure 5.3.
A reconstruction of the BQM-34 drone using an order 4 moment set. Even after 250 iterations, there is no perceptible sign of the basic drone shape. Additionally, the pixel error rate increases steadily with iterative attempts at improvement.

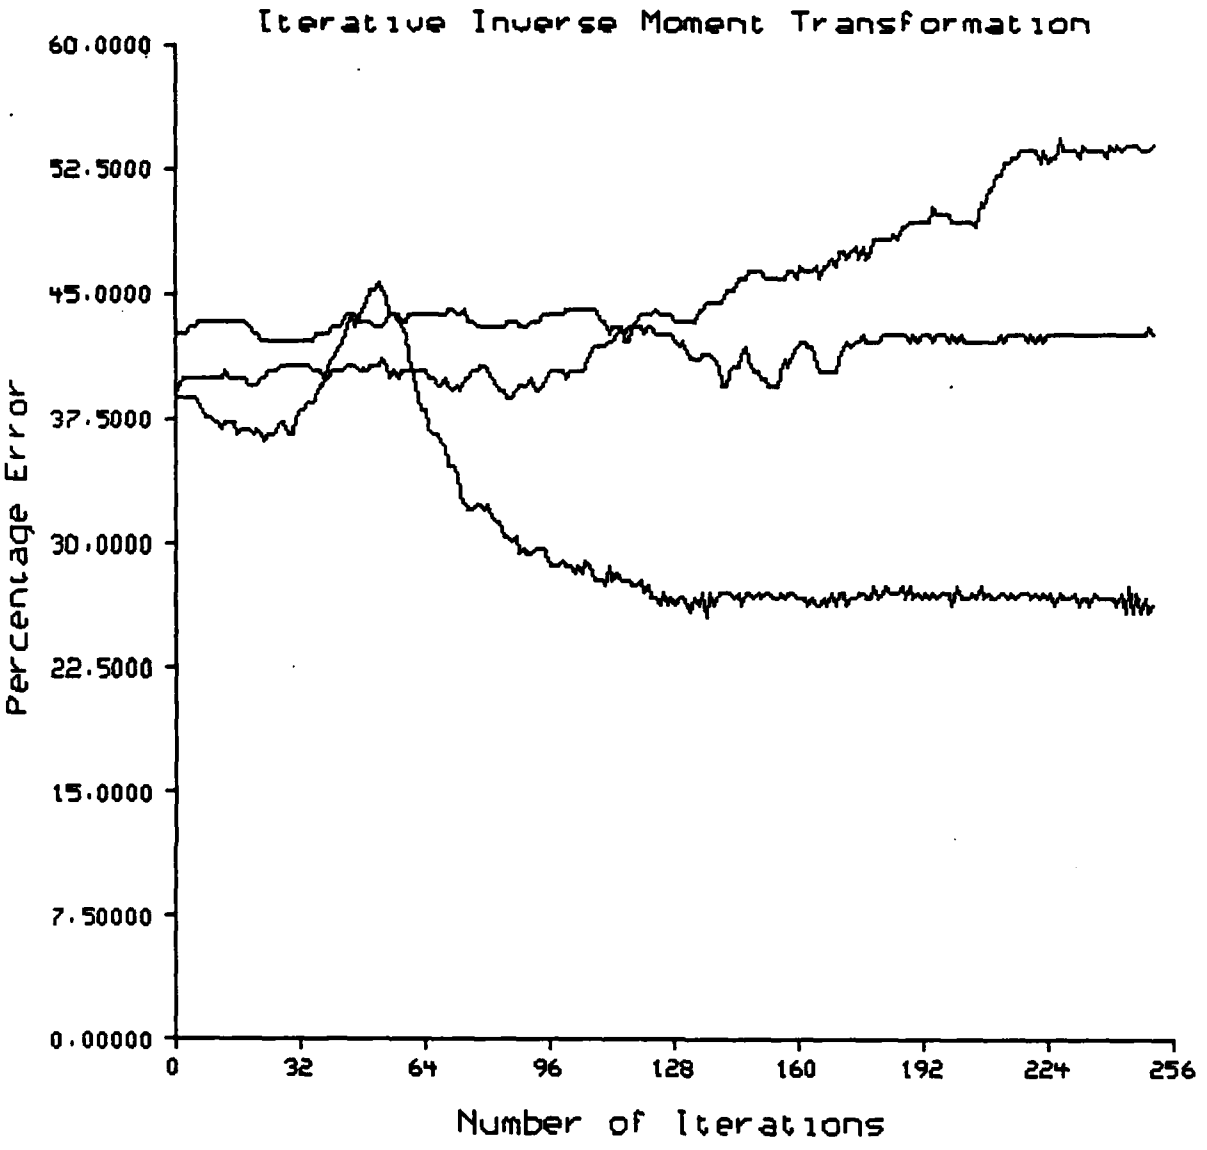

Figure 5.4.
A graph of the pixel based percentage of error versus the number of iterative attempts at reconstruction. At the 250th iteration, the top curve is for an order 4 moment set, the middle curve is for an order 6 moment set, and the bottom curve represents the results for an order 12 moment set iterative reconstruction.

Figure 5.5.
The results for an order 12 moment set reconstruction of a small rectangle. Note the quick convergence of the iterations to the basic shape with a feedback of 25%.

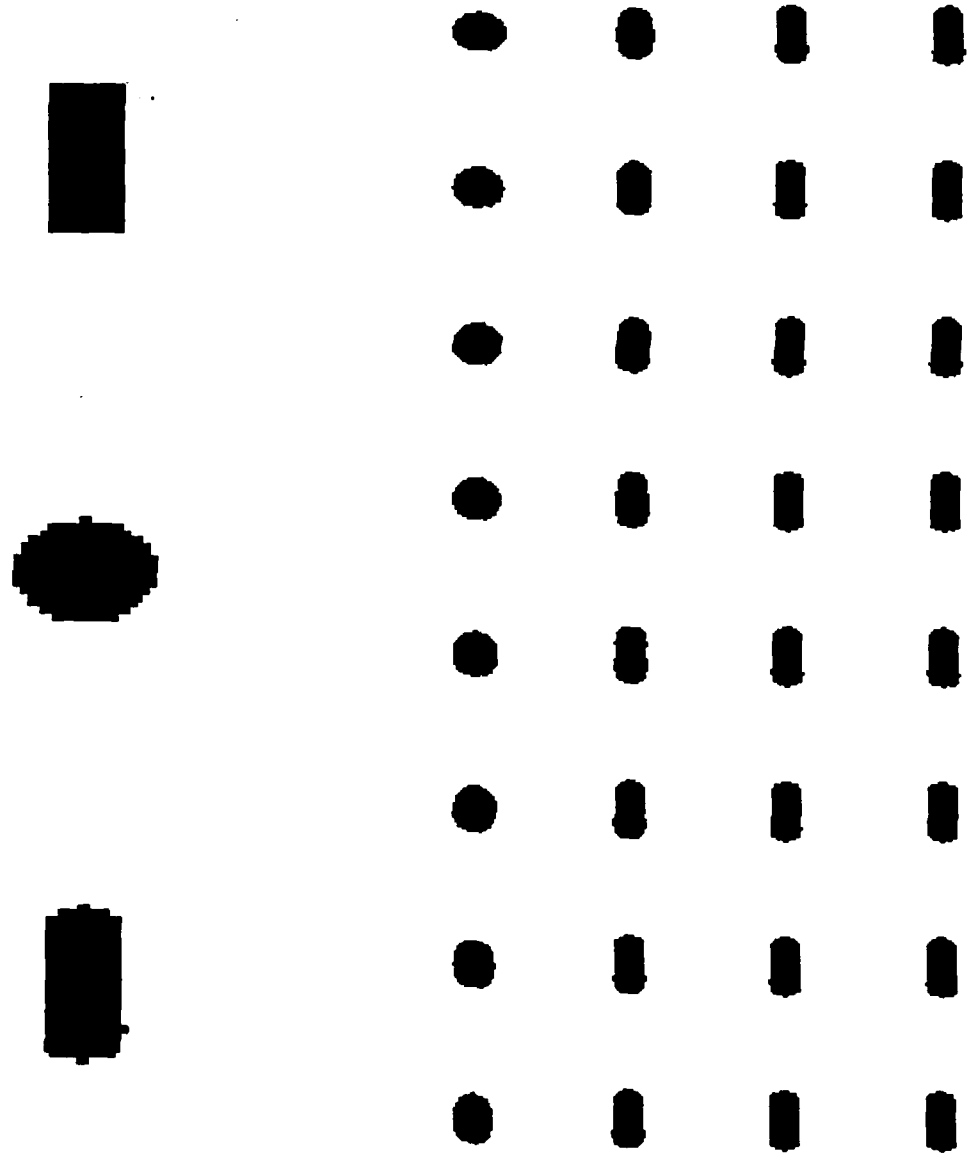

Figure 5.6.
An order 6 moment set reconstruction of the small rectangle.

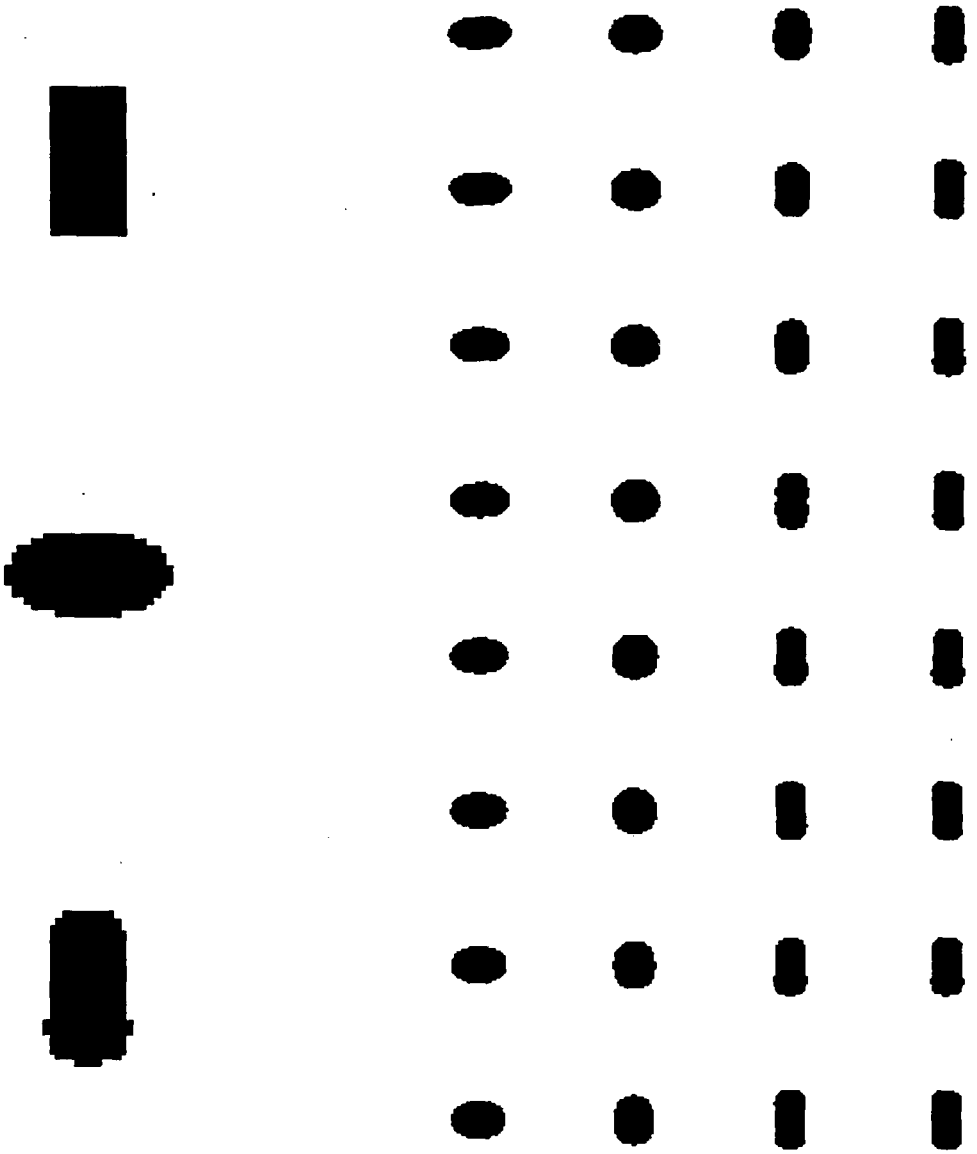

Figure 5.7.
Reconstruction of the small rectangle from a moment set of order 4. The results after 20 iterations are very comparable for the three different moment set order reconstructions of the small rectangle, all showing the basic shape and aspect ratio.

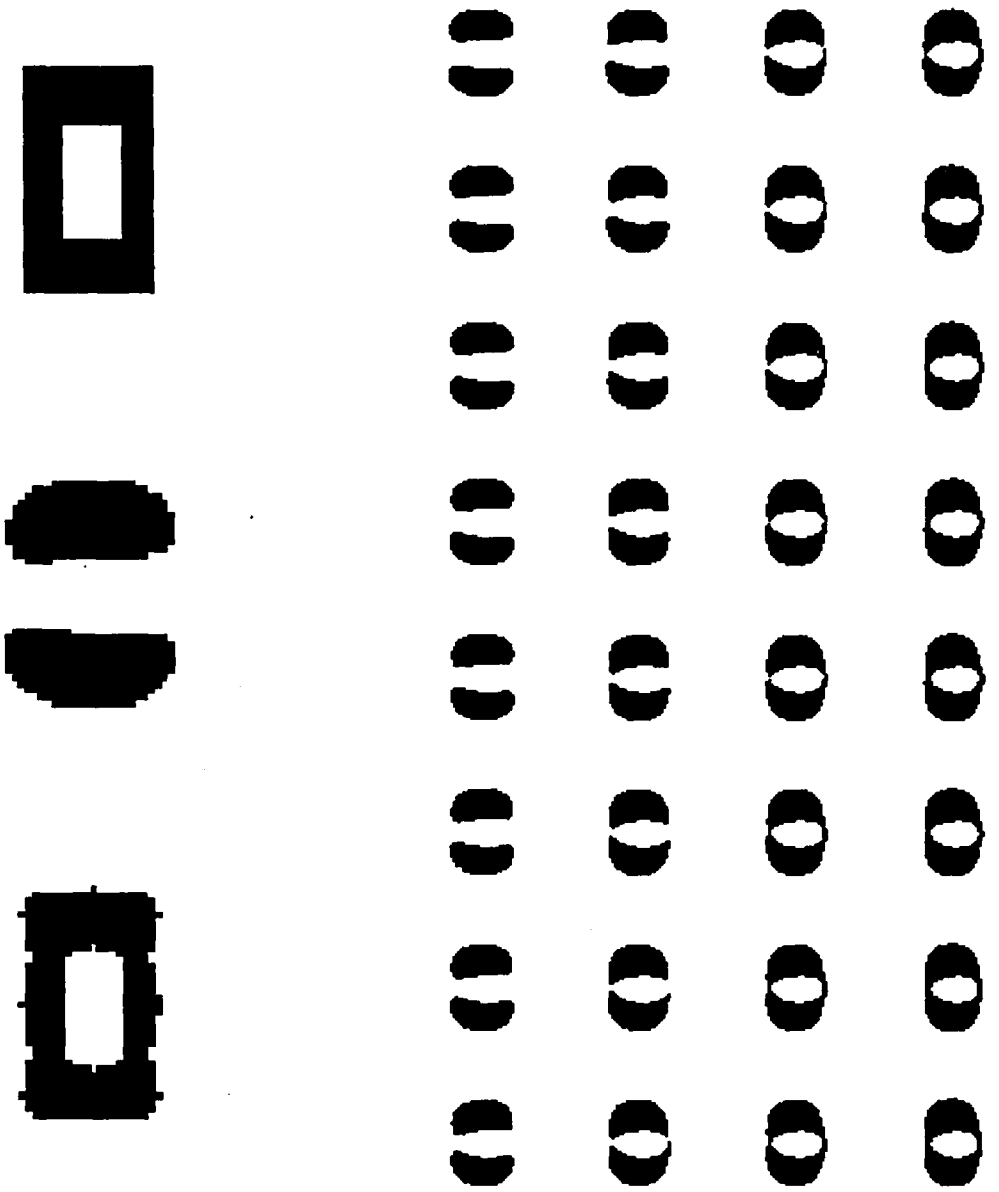

Figure 5.8.
The reconstruction of a more complex geometric shape. This hollow rectangle was produced from an order 12 moment set. An order of magnitude less feedback than for the small rectangle was used for this reconstruction. Higher feedback produced much faster shape change in the iterations, but the final result after 250 iterations was inferior to that shown.

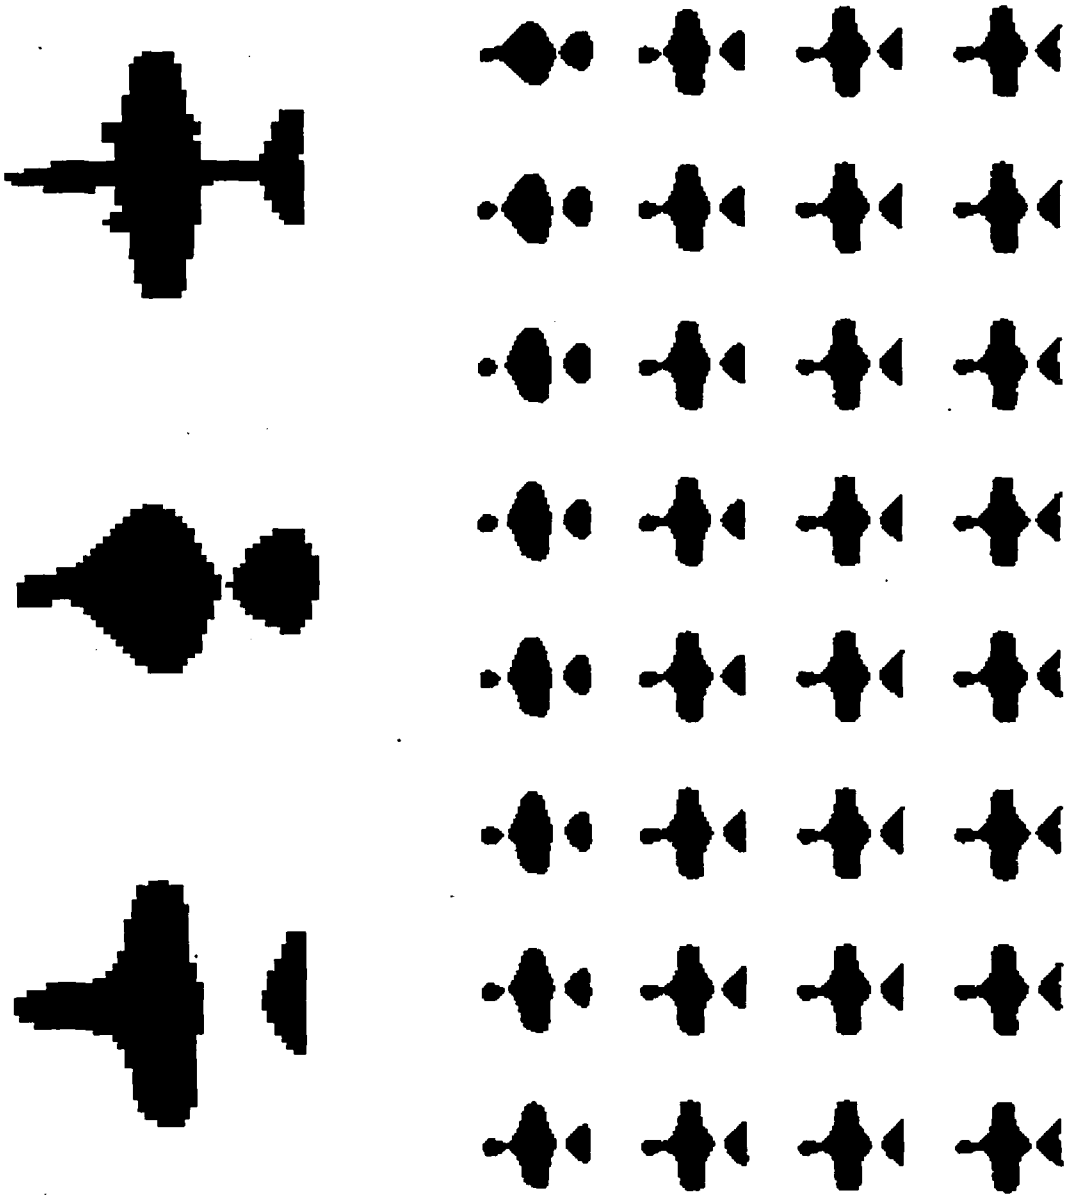

Figure 5.9.
An order 12 moment set reconstruction of the MiG-type aircraft. The shape is simpler than the drone example, which seems to produce improved results. The MiG also contains more object pixels than the drone, which allows for more effective image perturbation by the iterative technique. It is significant to note that even Teague's initial try with this image produces a recognizable result.

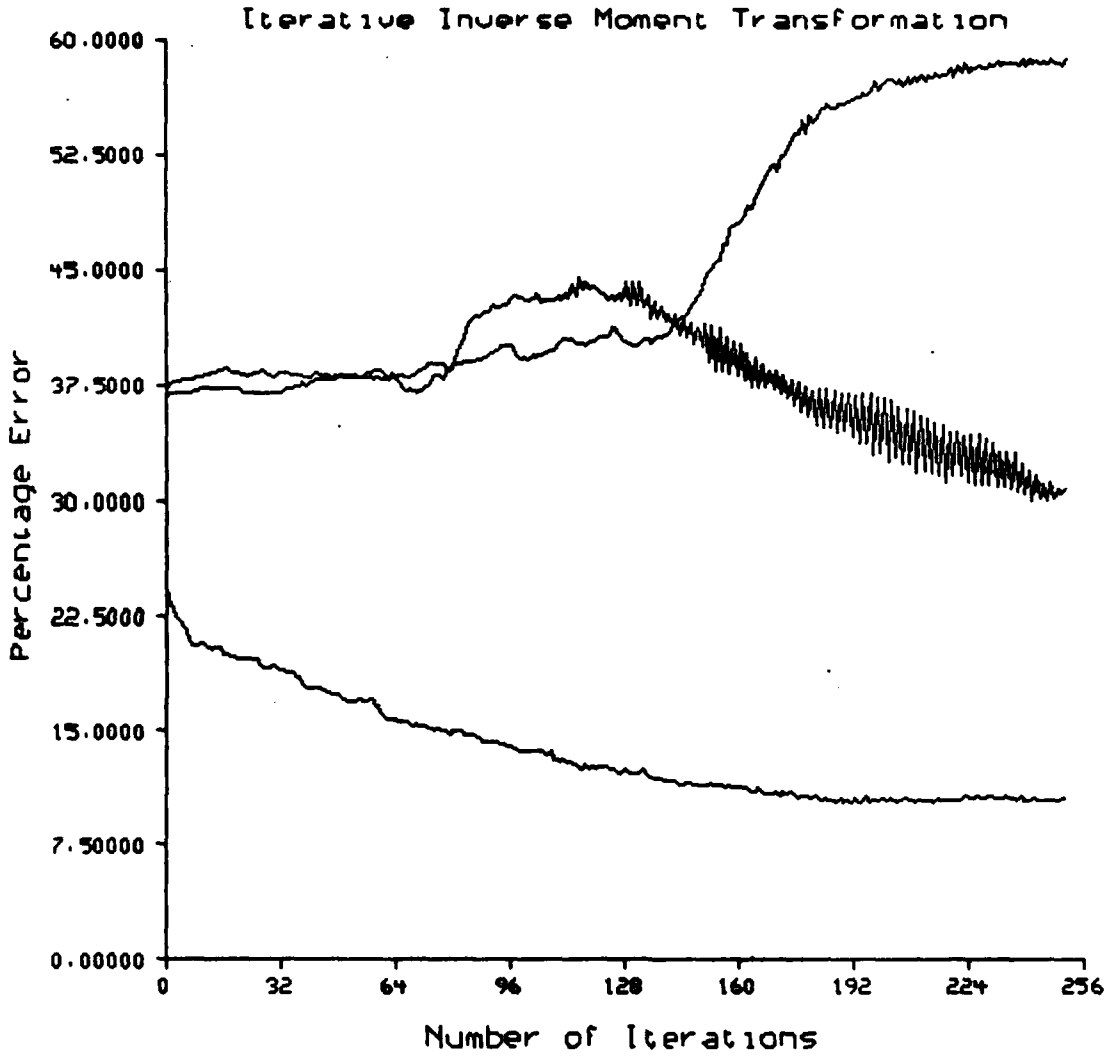

Figure 5.10.
This graph shows the results of iterative reconstructions of the MiG from moment sets of order 4, 6, and 12. These correspond to, respectively, the top, middle, and bottom curves at iteration 250. Note that the order 6 moment set reconstruction does not show improvement until after about 180 iterations. Fine tuning of the feedback to optimize results for a given order moment set might improve these results.

$\beta = 0.25$, the hollow rectangle used $\beta = 0.05$, and the BQM-34 had $\beta = 0.02$. The MiG used $\beta = 0.25$ for figure 5.9, and $\beta = 0.02$ for the graph in figure 5.10. As can be seen, the more complex and convoluted the shape, the lower the feedback used for stable iterations. Similarly, the feedback must be reduced, to maintain stability, as the order of the moment set used in the reconstrucion is decreased.

An additional operation was added to the basic algorithm that enhanced the results, particularly for the rectangular shapes. The procedure was to, effectively, zero a border region along the perimeter of the image in each reconstruction attempt. This implies some *a priori* knowledge of the object's dimensions, but for a small border, this is not too radical an approach. Removing spurious reconstructed areas along the image's edges was found to improve the algorithm's behavior.

A further modification to the original algorithm also allowed a feedback decay. The decay term is a value between 0 and 1 that the feedback is multiplied by after each iteration. thus the original algorithm can be modified by:

(6a) Reduce feedback as $\beta = \alpha^* \beta$, where α is a constant decay term

The decay factor can be used to reduce a potentially unstable feedback level before it begins to oscillate. This allows fast change for early iterations, and forces eventual convergence to a stable image. It should be noted that the running sum maintained in *image(i)* has the effect of reducing the effect of each succeeding iteration's feedback image, and thus there is an implicit feedback decay in the algorithm. The decay was $\alpha = 1.0$ for all the results presented here, in an effort to clearly display how the basic approach works, without the added obfuscation of the decay coefficient.

Tentative attempts at image reconstruction using a moment set of order 3 (10 numbers) indicated that it was virtually impossible to obtain stable iterations. For this reason, order 4 was selected as the lowest suitable order for our reconstruction attempts. This points out the correlation between stability of the iterative reconstruction approach, and the size of the com-

plete moment set used.

5.3. Conclusions on Image Reconstruction from Moments

The basic approach suggested by Teague [4] is inadequate to deal with silhouette imagery. The iterative feedback scheme developed here yields substantially improved results over the basic method. The complexity of the original silhouette dictates the success of the technique, as does the order of the complete moment set utilized in the reconstruction attempt.

The results for simple geometric shapes indicate that moment sets as small as order 4 may produce good reconstructions. More complex figures require a larger moment set, and some objects are difficult to reconstruct even with a moment set of order 12. The iterative technique also may tend to diverge from the original image for low order moment sets.

The number of pixels in the original object has a bearing on the success of the reconstruction. The fewer the pixels in the original object, the greater the sensitivity of the moments to slight pixel changes. Thus it is harder to produce stable iterations for a small object than for a large. This effect may, in some cases, explain why simpler geometric shapes are more accurately reconstructed than more complex, convoluted shapes. Part of the problem, though, is that sharp, angular silhouette changes are difficult to simulate with relatively low order polynomial approximations.

For the specific object examples reported here, the discussed conclusions are unilaterally verified. All objects showed their best results for order 12 moment set reconstructions, and then reduced performance for lower order attempts. Complex geometric objects required lower feedback than their simpler counterparts, to produce stable iterations that would result in an improved image reconstruction. And finally, the use of the ancilliary technique of border zeroing in some cases significantly enhances the results.

A reconstructive method may prove very important for gaining insight into moment classification ambiguities, since it gives an indication of both the significant features of the image being classified, and the sensitivity of the reconstruction to moment value variation.

Such variations might arise as artifacts of the sampling process at the image aquisition stage.

6. Overview and Conclusions

This report has considered various aspect of silhouette moments. The common thread concerns their effectiveness for the identification of three dimensional objects. At one extreme we have considered their suitability for classification of airplane silhouettes, and at the other we examined reconstructing an image from the very limited data set provided by its moments.

For the multiview airplane classification task reported in [8], a standard moment set with the specialized normalizations for aspect ratio and elemental variance showed performance far superior to Fourier descriptors and more *classical* moment invariant based approaches. These results held for both noisy and clean imagery. Furthermore, most objects could be identified with an order 3 moment set, and only a small increase in success was seen for the data set of section 2 when a higher order moment set was used.

The exhaustive test set of section 4 showed improvement when higher order moment sets were used. However, the best results occurred for order 5 moment sets, and only a slight increase in success was observed over order 3. The results for airplane identification using various subsets of the standard moment set indicated that odd order moments were more important than even order moments for classification purposes. A complete moment set of a order 5 still gave the best results, but a significant number of moment values could be eliminated by removing the order 4 moments, with only a slight degradation in classifier performance.

Section 3 considered noise reduction by image averaging. Both real FLIR data and synthetic data were used in the experimentation. The results indicate that, in situations where objects must be relocated within the image to achieve alignment for averaging purposes, averaging in the moment domain could produce better results than greylevel averaging in the image domain. Even in situations where the moment averaging provided similar, or even slightly inferior results, the tremendous computational simplification of object relocation in the moment domain still makes the technique attractive.

The final topic considered was silhouette reconstruction from a moment set. For this purpose Legendre moments were used, their orthogonal properties proving beneficial. An iterative reconstruction scheme using error feedback was developed, and results were shown for several images of varying size and complexity. It was found that reconstruction success was determined by the level of error feedback used, the order of the complete moment set used, and the composition of the original image. For some objects, such as a simple rectangle, the method was very successful, and for others, such as the BQM-34 drone model, little success was realized regardless of the parameters.

The ease of calculation of the moments of an image of an object, and the mathematical tractability of many operations such as scaling and translation in the moment domain commend the use of moments for object identification tasks. The results presented in this report further endorse the use of moments by showing the potential accuracy of object identifications from very small moment-based feature vectors.

7. References

1. T. P. Wallace and P. Wintz, "An Efficient, Three-Dimensional Aircraft Recognition Algorithm using Normalized Fourier Descriptors," *Computer Graphics and Image Processing* 3 pp. 99-126 (1980).
2. A. P. Reeves, "The General Theory of Moments for Shape Analysis and The Parallel Implementation of Moment Operations," Technical Report TR-EE 81-37 (October, 1981).
3. S. A. Dudani, K. J. Breeding, and R. B. McGhee, "Aircraft Identification by Moment Invariants," *IEEE Transactions on Computers* C-26(1) pp. 39-46 (January, 1977).
4. M. R. Teague,, "Image Analysis via the General Theory of Moments," *Journal of the Optical Society of America* 70(8) pp. 920-930 (August 1980).
5. A. P. Reeves, R. J. Prokop, S. E. Andrews, and F. P. Kuhl, "Three Dimensional Shape Analysis using Moments and Fourier Descriptors," *Proceedings of the Seventh International Conference on Pattern Recognition*, pp. 447-450 (July, 1984).
6. M-K Hu,, "Visual Pattern Recognition by Moment Invariants," *IRE Transactions on Information Theory* IT8 pp. 179-187 (Feb. 1962).
7. A. P. Reeves and A. Rostampour, "Shape Analysis of Segmented Objects Using Moments," *1981 Pattern Recognition and Image Processing Conference*, pp. 171-174 , Dallas, Texas(1981).
8. O. R. Mitchell, A. P. Reeves, and T. A. Grogan, "Algorithms and Architectures for Global Shape Analysis in Time-Varying Imagery," *SPIE Proceedings: Robotics and Industrial Inspection* 360 pp. 190-197 (August 1982).

Appendix A

THREE DIMENSIONAL SHAPE ANALYSIS USING MOMENTS AND FOURIER DESCRIPTORS

Anthony P. Reeves*, R. J. Prokop*, Susan E. Andrews*, Frank P. Kuhl**

* School of Electrical Engineering, Cornell University, Ithaca, NY 14853
** USA ARDC AMCCOM, Dover, NJ 07801

ABSTRACT

A procedure for using moment based feature vectors to identify a three-dimensional object from a two-dimensional image recorded at an arbitrary viewing angle and range is presented. Two global techniques which have been previously proposed for this task are Moments and Fourier descriptors. Nearly all the research in moments has centered on moment invariants; in this paper we are concerned with a different moment form called standard moments. We have used a standardized six airplane experiment in order to compare different techniques. Fourier descriptors, as used by Wallace and Wintz, and moment invariants have been compared with our own scheme for normalized moments. Various experiments have been conducted using mixtures of silhouette and boundary moments and different normalization techniques. Standard moments gave slightly better results than Fourier descriptors for this experiment; both of these techniques were much better than moment invariants.

INTRODUCTION

The properties of moments and Fourier descriptors have been examined by the use of an experiment involving six airplane shapes viewed at random angles in three dimensions. Fourier descriptors as used by Wallace and Wintz [1] have been compared with our own scheme [2] for normalized moments. Moment invariants as used by Dudani et. al. [3], have also been considered. Many different normalized moment forms have been proposed [2] including Legendre moments and variants of moment invariants [4]. In this paper the performances of several normalization techniques are compared.

Most of the previous work with moments, such as that by Dudani et al. [3], has centered on moment invariants; we have concentrated our efforts on normalized conventional moments which we consider to be less sensitive to noise. Normalization procedures for standard moments are described in [5]. A new normalization technique, aspect ratio normalization, is described in this paper. Fourier descriptors of the contour of a shape may be used in a similar way as moments to identify objects ; [6] the method used in our experiments is fully described in [1]. The moment invariants used in our experiment are similar to those used by Dudani except that the range to the object is not known.

The Six Airplane Experiment

A standardized six airplane classification task was used for all experiments. It is loosely based on an experiment designed and developed by Wallace and Wintz [1]. Their experiment was, in turn, loosely based on the experiment conducted by Dudani et.al. [3].

The basic six airplane experiment is as follows. The six airplanes are described by a short list of parameters. These parameters are used to generate a representation of the three dimensional surface of an airplane consisting of a set of three and four sided polygons. A library of feature vector sets of two dimensional views for each airplane is generated. Library samples are usually taken over a uniform grid of viewing angles. Each library entry is normalized with respect to size, location, and in-plane rotation.

A set of random views of the six airplanes is used to generate a test set of feature vectors. Each test view is matched by means of nearest Euclidian neighbor in feature space to its closest entry in the library. The class of this entry is used to identify the test airplane. The performance of the shape analysis scheme is indicated by the number of correct classifications.

MOMENT FEATURE VECTORS

The conventional definition of the two-dimensional moment of order (p+q) of a function f(x,y) is

$$M_{pq} = \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} x^p y^q f(x,y)\, dx\, dy \qquad p,q = 0,1,2, \ldots \qquad (1)$$

A set of moment values may be used to represent a segment of an image. Operations such as rotation, translation and scale change may more easily be achieved in the moment domain than the original pixel domain. Furthermore, a truncated set of moments may offer a more convenient and economical representation of the essential shape characteristics of an image segment than a pixel based representation.

Recent work [5] has shown that both the moments of the binary silhouette and the grey-valued image segment may be used together for classification. Furthermore, the two dimensional moments may be used to obtain moments of projections through the image segment which may, in turn, be used in classification schemes.

Presented at the 7th International Conference on Pattern Recognition, Montreal, July 1984.

Standard Moments

A set of standard moments is a CMS which has been normalized with respect to scale, translation, rotation and optionally reflection. The low order standard moments have the following values:

$M_{00} = 1$ (area normalized to 1)
$M_{10} = 0$ (central moments, taken about
$M_{01} = 0$ the center of gravity)

$M_{11} = 0$ (rotation normalization; moments are taken about the principal axes)

For rotation normalization we also have

$$M_{20} \geq M_{02}$$
$$M_{30} \geq 0$$

If reflection normalization is selected

$$M_{03} \geq 0$$

The standard moment representation of a segment is very convenient for shape analysis and may be compared directly with other standard moment sets for shape similarity.

Aspect Ratio Normalization

A feature of the standard moments, as described in the previous section, is that for blob-like objects the magnitude of the moments will decrease as the order increases. The condition for diminishing magnitudes is that the object be confined within a 2x2 square centered at the origin. Recall that M_{00}, the object area, is 1, therefore all other moments will be less than 1. If any part of the object extends beyond this square region, then some moments will become increasingly large and even for low orders may be much larger than 1. A significant number of the airplane library images extend beyond the 2x2 square.

In order to improve the behavior of feature vectors, an Aspect ratio normalization was developed. This is a projective transform of an image of the object which converts the ellipsoid of inertia of the object to a circle, i.e., $M_{20} = M_{02}$, while maintaining $M_{00} = 1$. There are still shapes which would extend beyond the 2x2 square even after this normalization; however, all of the airplane images were contained in this square.

A change in scale by a factor of δ in the x direction and γ in the y direction is given by:

$$M'_{pq} = \delta^{1+p} \gamma^{1+q} M_{pq} \qquad (2)$$

For aspect normalization the area remains 1, and $M_{20} = M_{02}$ therefore,

$$\delta = (M_{02}/M_{20})^{0.25}, \quad \gamma = 1/\delta \qquad (3)$$

This transformation is applied to standard moments. For the transformed moments, $M_{20} = M_{02}$, therefore we discard one of them. The aspect ratio δ is an important shape feature and is used as the first or most significant feature element. Its value is in the range $0 < \delta \leq 1$.

CLASSIFICATION PROCEDURES

Each airplane, considered from all possible viewing angles, forms a continuous surface in feature space. The novel classification problem is to find the closest airplane surface to an unknown airplane, which is represented by a single point in feature space. Classical feature vector conditioning procedures, such as the Eigen value transformation as used by Wallace and Wintz, are not necessarily appropriate in this case since the within-class feature vector variation (for all different possible viewing angles) is, in general, less than the between-class variation for a given viewing angle.

We have used two classification criterias. The first is simply minimum Euclidian distance in feature space. Both standard moments and normalized Fourier descriptors have feature elements which decrease in magnitude with order or frequency. Therefore, more weight is placed on the lower order or frequency elements which are expected to be more robust with respect to noise.

The second criteria is minimum Euclidian distance after "variance balancing" the feature vectors. Variance balancing is a weighting technique to make the effect of the feature elements more equal. The variance of each feature element for all feature vectors in the total library for all six airplanes is first computed. Then all feature elements are weighted by the inverse of the standard deviation computed for the library. Therefore, if all views of the airplane are equally likely, the expected variance of a weighted feature vector is one.

EXPERIMENTS AND RESULTS

The Basic Experiment

The following parameters were used in the basic experiment.

(1) The library views consisted of 500 views uniformly sampled over the total sphere of viewing angles. In Dudani's experiment 550 views of each airplane were taken at 5 degree intervals in a 140 by 90 degree sector. The library views used by Wallace and Wintz consisted of 143 entries covering one hemisphere in a 11 x 13 grid.

(2) The test set consisted of the same 50 random views of each airplane for a total set size of 300. The viewing angles used were identical to those of Wallace and Wintz.

(3) Library entries were computed from two-dimensional views digitized on a 128 x 128 pixel matrix. Test set entries were also computed from a 128 x 128 pixel representation.

Representative views of each of the six airplanes with a 128 x 128 pixel resolution are shown in Fig. 1. Typical random views of these airplanes are shown in Fig 2. A noisy version of the test set of airplanes was also used in the experiments.

Noisy data generation

Considerable care was taken in the generation of the noisy data to give realistic results. The following operations were applied to each clean test image.

1. The silhouette image was converted to a grey level image in which each object point is 64 greater than the background.

2. A 2 x 2 local mean filter operation is performed which blurs the edges of the grey level image.
3. Gaussian noise is added with a mean of zero and a standard deviation of 30.
4. A second 2 x 2 local mean filter operation is performed.
5. A thresholding operation is performed at the optimal grey level for the image specified in step 1.
6. A binary filtering algorithm is applied which selects only the main connected region of ones (object points) in the image.

Some examples of noisy images are given in Fig. 3.

Results

The classification results for Fourier descriptors are shown in Fig. 4. For small numbers of feature vector elements variance balancing is beneficial but for larger numbers of elements the results are worse. The results for clean data are slightly better than those obtained by Wallace and Wintz; this could be due to the more dense library which we used.

The classification results for conventional moments and aspect ratio normalized moments are shown in Fig. 5. Variance balancing makes a significant improvement in the results for clean data in both cases. The best results are obtained with variance balanced aspect ratio normalized moments. This performance is effectively achieved with moments of order 3; using more moment values has very little effect. Variance balancing is also very important for the noisy data. There is a significant improvement in adding more feature vector elements in this case.

Fig 6. shows the performance of various types of feature vectors on clean data using variance balancing. Combining silhouette and boundary data did give a slight improvement over the silhouette data alone. Moments based on Legendre polynomials have been proposed by Teague. While the orthogonal property might be expected to lead to better feature vector representation, our results indicate that they performed worse than the silhouette moments from which they were derived. Boundary moments by themselves performed much worse than silhouette moments.

Also shown in Fig 6. are results for moment invariants which are worse than any other technique tested. These were computed in a similar way to Dudani except for the first feature element value. Dudani assumed that the distance to the object was known; therefore the absolute scale of the object was also known and is embodied in the first feature vector element. In our studies we have assumed that the distance to the object is not known. In our moment invariants the radius of gyration (the first feature element) is size normalized by the perceived area of the object M_{00} while Dudani ignores this area moment value.

CONCLUSIONS

The following summarizes the best results which have been obtained for each feature vector type:

Feature Vector Length	% Correctly Classified: Standard Moments	Fourier Descriptors
6	93.0 (85.7)	77.7 (75.0)
11	93.3 (89.7)	90.7 (82.3)
17	93.3 (89.7)	90.3 (85.0)
24	92.7 (91.0)	91.7 (85.6)

The figures in parenthesis are for a noisy set of unknowns. For moment invariants of the silhouette (7 feature vectors) the highest recognition rate was 68.0% for clean data and 69.3% for noisy data. Adding boundary moment invariants gave worse results. The above results indicate that there is very little difference between the performance of Fourier descriptors and standard moments for this task; both do very well. Moment invariants did considerably worse than the other techniques. The method used for normalization significantly affects the results. In general, the addition of boundary moments did not improve the classification results for a given number of feature vectors and boundary moments by themselves were not very good.

Acknowledgements

The contributions made by A. Rostampour and C. Martin to early stages of this work are gratefully acknowledged. This work was funded in part by the National Science Foundation Grant No ECS-8303583 and the Army Research Office.

REFERENCES

1. T. P. Wallace and P. Wintz, "An Efficient, Three-Dimensional Aircraft Recognition Algorithm using Normalized Fourier Descriptors," Computer Graphics and Image Processing Vol. 3 pp. 99-126 (1980).
2. A. P. Reeves, "The General Theory of Moments for Shape Analysis and The Parallel Implementation of Moment Operations," Technical Report TR-EE 81-37 (October, 1981).
3. S. A. Dudani, K. J. Breeding, and R. B. McGhee, "Aircraft Identification by Moment Invariants," IEEE Transactions on Computers Vol. C-26(1) pp. 39-46 (January, 1977).
4. M. R. Teague., "Image Analysis via the General Theory of Moments," Journal of the Optical Society of America Vol. 70(8) pp. 920-930 (August 1980).
5. A. P. Reeves and A. Rostampour, "Shape Analysis of Segmented Objects Using Moments," 1981 Pattern Recognition and Image Processing Conference, pp. 171-174 , Dallas, Texas(1981).
6. O. R. Mitchell, A. P. Reeves, and T. A. Grogan, "Algorithms and Architectures for Global Shape Analysis in Time-Varying Imagery," SPIE Proceedings: Robotics and Industrial Inspection Vol. 360 pp. 190-197 (August 1982).

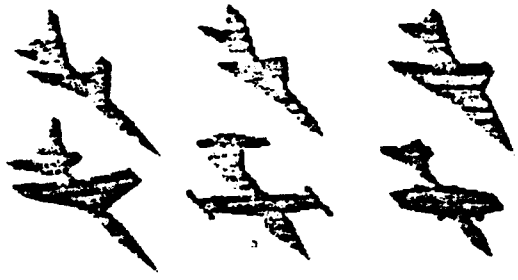

Fig 1. The Six Airplanes

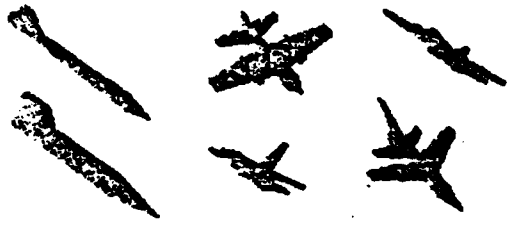

Fig 2. Examples of Clean Images

Fig 3. Examples of Noisy Images

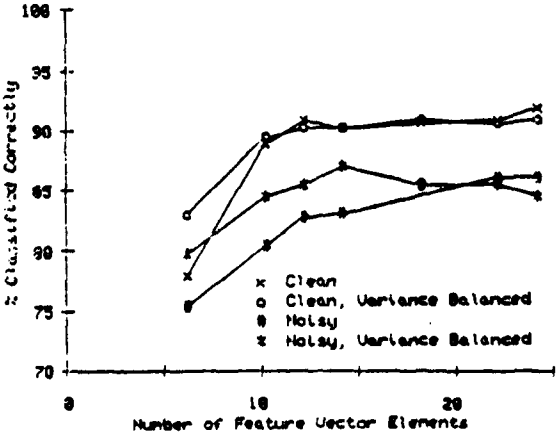

Fig 4. Classification results for Normalized Fourier Descriptors

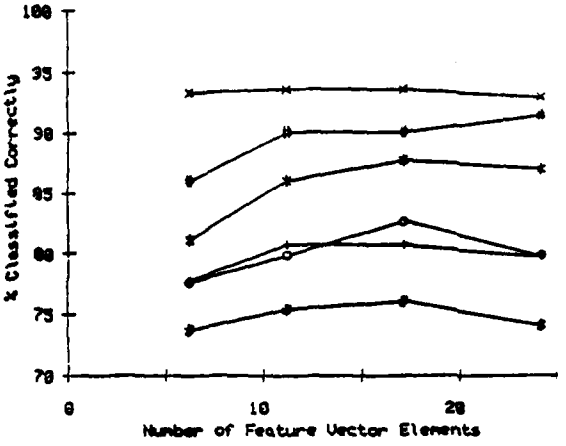

Fig 5. Classification results for Silhouette Moments

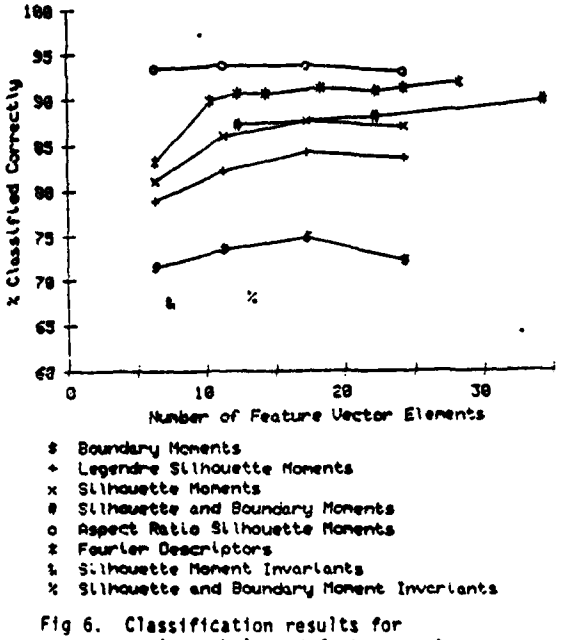

Fig 6. Classification results for variance balanced feature vectors

END

FILMED

7-85

DTIC